

---

# Leveraging Good Representations in Linear Contextual Bandits

---

Matteo Papini<sup>\*†1</sup> Andrea Tirinzoni<sup>\*1</sup> Marcello Restelli<sup>1</sup> Alessandro Lazaric<sup>2</sup> Matteo Pirotta<sup>2</sup>

## Abstract

The linear contextual bandit literature is mostly focused on the design of efficient learning algorithms for a given representation. However, a contextual bandit problem may admit multiple linear representations, each one with different characteristics that directly impact the regret of the learning algorithm. In particular, recent works showed that there exist “good” representations for which constant problem-dependent regret can be achieved. In this paper, we first provide a systematic analysis of the different definitions of “good” representations proposed in the literature. We then propose a novel selection algorithm able to adapt to the best representation in a set of  $M$  candidates. We show that the regret is indeed never worse than the regret obtained by running LINUCB on the best representation (up to a  $\ln M$  factor). As a result, our algorithm achieves constant regret whenever a “good” representation is available in the set. Furthermore, we show that the algorithm may still achieve constant regret by implicitly constructing a “good” representation, even when none of the initial representations is “good”. Finally, we empirically validate our theoretical findings in a number of standard contextual bandit problems.

## 1. Introduction

The stochastic contextual bandit is a general framework to formalize sequential decision-making problems in which at each step the learner observes a context drawn from a fixed distribution, it plays an action, and it receives a noisy reward. The goal of the learner is to maximize the reward accumulated over  $n$  rounds, and the performance is typically measured by the regret w.r.t. playing the optimal action in each context. This paradigm has found application in a large

---

<sup>\*</sup>Equal contribution <sup>†</sup>Work done while at Facebook AI Research <sup>1</sup>Politecnico di Milano, Milan, Italy <sup>2</sup>Facebook AI Research, Paris, France. Correspondence to: Matteo Papini <matteo.papini@polimi.it>.

range of domains, including recommendation systems, on-line advertising, and clinical trials (e.g., Bouneffouf & Rish, 2019). Linear contextual bandit (Lattimore & Szepesvári, 2020) is one of the most studied instances of contextual bandit due to its efficiency and strong theoretical guarantees. In this setting, the reward for each context  $x$  and action  $a$  is assumed to be representable as the linear combination between  $d$ -dimensional features  $\phi(x, a) \in \mathbb{R}^d$  and an unknown parameter  $\theta^* \in \mathbb{R}^d$ . In this case, we refer to  $\phi$  as a realizable representation. Algorithms based on the optimism-in-the-face-of-uncertainty principle such as LINUCB (Chu et al., 2011) and OFUL (Abbasi-Yadkori et al., 2011), have been proved to achieve minimax regret bound  $O(Sd\sqrt{n} \ln(nL))$  and problem-dependent regret  $O(\frac{S^2 d^2}{\Delta} \ln^2(nL))$ , where  $\Delta$  is the minimum gap between the reward of the best and second-best action across contexts, and  $L$  and  $S$  are upper bounds to the  $\ell_2$ -norm of the features  $\phi$  and  $\theta^*$ , respectively.

Unfortunately, the dimension  $d$ , and the norm upper bounds  $L$  and  $S$ , are not the only characteristics of a representation to have an effect on the regret and existing bounds may fail at capturing the impact of the context-action features on the performance of the algorithm. In fact, as illustrated in Fig. 1, running LINUCB with different realizable representations with same parameters  $d$  and  $S$  may lead to significantly different performance. Notably, there are “good” representations for which LINUCB achieves *constant* regret, i.e., not scaling with the horizon  $n$ . Recent works identified different conditions on the representation that can be exploited to achieve constant regret for LINUCB (Hao et al., 2020; Wu et al., 2020). Similar conditions have also been leveraged to prove other interesting learning properties, such as sub-linear regret for greedy algorithms (Bastani et al., 2020), or regret guarantees for model selection between linear and multi-arm representations (Chatterji et al., 2020; Ghosh et al., 2020). While all these conditions, often referred to as *diversity conditions*, depend on how certain context-arm features span the full  $\mathbb{R}^d$  space, there is no systematic analysis of their connections and of which ones can be leveraged to achieve constant regret in linear contextual bandits.

In this paper, we further investigate the concept of “good” representations in linear bandit and we provide the following contributions: **1)** We review the diversity conditions available in the literature, clarify their relationships, and discuss how they are used. We then focus on our primary




Figure 1. Regret of LINUCB with different *realizable* representations with same dimension  $d$  and parameter bound  $S$ . The dashed blue line is LEADER, our proposed representation selection algorithm. Details in App. G.1.

goal, which is to characterize the assumptions needed to achieve constant regret for LINUCB. **2)** We introduce a novel algorithm that effectively selects the best representation in a given set, thus achieving constant regret whenever at least one “good” representation is provided. **3)** Furthermore, we show that, in certain problems, the algorithm is able to combine given representations to implicitly form a “good” one, thus achieving constant problem-dependent regret even when running LINUCB on any of the representations would not. **4)** Finally, we empirically validate our theoretical findings on a number of contextual bandit problems.

**Related work.** The problem of selecting the best representation in a given set can be seen as a specific instance of the problem of *model selection* in bandits. In model selection, the objective is to choose the best candidate in a set of *base learning algorithms*. At each step, a *master algorithm* is responsible for selecting a base algorithm, which in turn prescribes the action to play and the reward is then provided as feedback to the base algorithms. Examples of model selection methods include adversarial masters –e.g., EXP4 (Auer et al., 2002; Maillard & Munos, 2011) and CORRAL (Agarwal et al., 2017; Pacchiano et al., 2020b)– and stochastic masters (Abbasi-Yadkori et al., 2020; Lee et al., 2020; Bibaut et al., 2020; Pacchiano et al., 2020a). For a broader discussion refer to App. A or (Pacchiano et al., 2020a, Sec. 2). Most of these algorithms achieve the regret of the best base algorithm up to a polynomial dependence on the number  $M$  of base algorithms (Agarwal et al., 2017). While existing model selection methods are general and can be applied to any type of base algorithms,<sup>1</sup> they may not be effective in problems with a specific structure.

An alternative approach is to design the master algorithm for a specific category of base algorithms. An instance of

<sup>1</sup>Most of existing methods only require prior knowledge of the regret of the optimal base algorithm or a bound on the regret of all base algorithms. CORRAL also requires the base algorithms to satisfy certain stability conditions.

this case is the representation-selection problem, where the base algorithms only differ by the representation used to estimate the reward. Foster et al. (2019) and Ghosh et al. (2020) consider a set of nested representations, where the best representation is the one with the smallest dimensionality for which the reward is realizable. Finally, Chatterji et al. (2020) focus on the problem of selecting between a linear and a multi-armed bandit representation. In this paper, we consider an alternative representation-selection problem in linear contextual bandits, where the objective is to exploit constant-regret “good” representations. Differently from our work, Lattimore et al. (2020) say that a linear representation is “good” if it has a low *misspecification* (i.e., it represents the reward up to a small approximation error), while we focus on *realizable* representations for which LINUCB achieves constant-regret.

## 2. Preliminaries

We consider the stochastic contextual bandit problem (*contextual problem* for short) with context space  $\mathcal{X}$  and finite action set  $\mathcal{A} = [K] = \{1, \dots, K\}$ . At each round  $t \geq 1$ , the learner observes a context  $x_t$  sampled i.i.d. from a distribution  $\rho$  over  $\mathcal{X}$ , it selects an arm  $a_t \in [K]$  and it receives a reward  $y_t = \mu(x_t, a_t) + \eta_t$  where  $\eta_t$  is a  $\sigma$ -subgaussian noise. The learner’s objective is to minimize the pseudo-regret  $R_n = \sum_{t=1}^n \mu^*(x_t) - \mu(x_t, a_t)$  for any  $n > 0$ , where  $\mu^*(x_t) := \max_{a \in [K]} \mu(x_t, a)$ . We define the minimum gap as  $\Delta = \inf_{x \in \mathcal{X}: \rho(x) > 0, a \in [K], \Delta(x, a) > 0} \{\Delta(x, a)\}$  where  $\Delta(x, a) = \mu^*(x) - \mu(x, a)$ . A *realizable*  $d_\phi$ -dimensional linear representation is a feature map  $\phi : \mathcal{X} \times [K] \rightarrow \mathbb{R}^{d_\phi}$  for which there exists an unknown parameter vector  $\theta_\phi^* \in \mathbb{R}^{d_\phi}$  such that  $\mu(x, a) = \langle \phi(x, a), \theta_\phi^* \rangle$ . When a realizable linear representation is available, the problem is called (stochastic) linear contextual bandit and can be solved using, among others, optimistic algorithms like LINUCB (Chu et al., 2011) or OFUL (Abbasi-Yadkori et al., 2011).

Given a realizable representation  $\phi$ , at each round  $t$ , LINUCB builds an estimate  $\theta_{t\phi}$  of  $\theta_\phi^*$  by ridge regression using the observed data. Denote by  $V_{t\phi} = \lambda I_{d_\phi} + \sum_{k=1}^{t-1} \phi(x_k, a_k) \phi(x_k, a_k)^\top$  the  $(\lambda > 0)$ -regularized design matrix at round  $t$ , then  $\theta_{t\phi} = V_{t\phi}^{-1} \sum_{k=1}^{t-1} \phi(x_k, a_k) y_k$ . Assuming that  $\|\theta_\phi^*\|_2 \leq S_\phi$  and  $\sup_{x, a} \|\phi(x, a)\|_2 \leq L_\phi$ , LINUCB builds a confidence ellipsoid  $\mathcal{C}_{t\phi}(\delta) = \{\theta \in \mathbb{R}^{d_\phi} : \|\theta_{t\phi} - \theta\|_{V_{t\phi}} \leq \beta_{t\phi}(\delta)\}$ . As shown in (Abbasi-Yadkori et al., 2011, Thm. 1), when

$$\beta_{t\phi}(\delta) := \sigma \sqrt{2 \ln \left( \frac{\det(V_{t\phi})^{1/2} \det(\lambda I_{d_\phi})^{-1/2}}{\delta} \right)} + \sqrt{\lambda} S_\phi,$$

then  $\mathbb{P}(\forall t \geq 1, \theta_\phi^* \in \mathcal{C}_{t\phi}(\delta)) \geq 1 - \delta$ . At each step  $t$ , LINUCB plays the action with the highest upper-confidence bound  $a_t = \operatorname{argmax}_{a \in [K]} \max_{\theta \in \mathcal{C}_{t\phi}(\delta)} \langle \phi(x_t, a), \theta \rangle$ , and it is shown to achieve a regret bounded as reported in the

following proposition.

**Proposition 1** (Abbasi-Yadkori et al., 2011, Thm. 3, 4). *For any linear contextual bandit problem with  $d_\phi$ -dimensional features,  $\sup_{x,a} \|\phi(x, a)\|_2 \leq L_\phi$ , an unknown parameter vector  $\|\theta_\phi^*\|_2 \leq S_\phi$ , with probability at least  $1 - \delta$ , LINUCB suffers regret  $R_n = O(S_\phi d_\phi \sqrt{n} \ln(nL_\phi/\delta))$ . Furthermore, if the problem has a minimum gap  $\Delta > 0$ , then the regret is bounded as<sup>2</sup>  $R_n = O\left(\frac{S_\phi^2 d_\phi^2}{\Delta} \ln^2(nL_\phi/\delta)\right)$ .*

In the rest of the paper, we assume w.l.o.g. that all terms  $\lambda$ ,  $\Delta_{\max} = \max_{x,a} \Delta(x, a)$ ,  $S_\phi$ ,  $\sigma$  are larger than 1 to simplify the expression of the bounds.

### 3. Diversity Conditions

Several assumptions, usually referred to as *diversity conditions*, have been proposed to define linear bandit problems with specific properties that can be leveraged to derive improved learning results. While only a few of them were actually leveraged to derive constant regret guarantees for LINUCB (others have been used to prove e.g., sub-linear regret for the greedy algorithm, or regret guarantees for model selection algorithms), they all rely on very similar conditions on how certain context-action features span the full  $\mathbb{R}^{d_\phi}$  space. In this section, we provide a thorough review of these assumptions, their connections, and how they are used in the literature. As diversity conditions are getting more widely used in bandit literature, we believe this review may be of independent interest. Sect. 4 will then specifically focus on the notion of *good* representation for LINUCB.

We first introduce additional notation. For a realizable representation  $\phi$ , let  $\phi^*(x) := \phi(x, a_x^*)$ , where  $a_x^* \in \operatorname{argmax}_{a \in [K]} \mu(x, a)$  is an optimal action, be the vector of *optimal features* for context  $x$ . In the following we make the assumption that  $\phi^*(x)$  is unique. Also, let  $\mathcal{X}^*(a) = \{x \in \mathcal{X} : \mu(x, a) = \mu^*(x)\}$  denote the set of contexts where  $a$  is optimal. Finally, for any matrix  $A$ , we denote by  $\lambda_{\min}(A)$  its minimum eigenvalue. For any contextual problem with reward  $\mu$  and context distribution  $\rho$ , the diversity conditions introduced in the literature are summarized in Tab. 2 together with how they were leveraged to obtain regret bounds in different settings.<sup>3</sup>

We first notice that all conditions refer to the smallest eigenvalue of a design matrix constructed on specific context-action features. In other words, diversity conditions require certain features to span the full  $\mathbb{R}^{d_\phi}$  space. The

<sup>2</sup>The logarithmic bound reported in Prop. 1 is slightly different than the one in (Abbasi-Yadkori et al., 2011) since we do not assume that the optimal feature is unique.

<sup>3</sup>In some cases, we adapted conditions originally defined in the disjoint-parameter setting, where features only depend on the context (i.e.,  $\phi(x)$ ) and the unknown parameter  $\theta_a^*$  is different for each action  $a$ , to the shared-parameter setting (i.e., where features are functions of both contexts and actions) introduced in Sect. 2.

non-redundancy condition is a common technical assumption (e.g., Foster et al., 2019) and it simply defines a problem whose dimensionality cannot be reduced without losing information. Assuming the context distribution  $\rho$  is full support, BBK and CMB are structural properties of the representation that are independent from the reward. For example, BBK requires that, for each action, there must be feature vectors lying in all orthants of  $\mathbb{R}^{d_\phi}$ . In the case of finite contexts, this implies there must be at least  $2^{d_\phi}$  contexts. WYS and HLS involve the notion of reward optimality. In particular, WYS requires that all actions are optimal for at least a context (in the continuous case, for a non-negligible set of contexts), while HLS only focuses on optimal actions.

We now review how these conditions (or variations thereof) were applied in the literature. CMB is a rather strong condition that requires the features associated with each individual action to span the whole  $\mathbb{R}^{d_\phi}$  space. Chatterji et al. (2020) leverage a CMB-like assumption to prove regret bounds for OSOM, a model-selection algorithm that unifies multi-armed and linear contextual bandits. More precisely, they consider a variation of CMB, where the context distribution induces stochastic feature vectors for each action that are independent and centered. The same condition was adopted by Ghosh et al. (2020) to study representation-selection problems and derive algorithms able to adapt to the (unknown) norm of  $\theta_\phi^*$  or select the smallest realizable representation in a set of nested representations. Bastani et al. (2020, Assumption 3) introduced a condition similar to BBK for the *disjoint-parameter* case. In their setting, they prove that a non-explorative greedy algorithm achieves  $O(\ln(n))$  problem-dependent regret in linear contextual bandits (with 2 actions).<sup>4</sup> Hao et al. (2020, Theorem 3.9) showed that HLS representations can be leveraged to prove constant problem-dependent regret for LINUCB in the shared-parameter case. Concurrently, Wu et al. (2020) showed that, under WYS, LINUCB achieves constant expected regret in the disjoint-parameter case. A WYS-like condition was also used by Bastani et al. (2020, Assumption 4) to extend the result of sublinear regret for the greedy algorithm to more than two actions. The relationship between all these conditions is derived in the following lemma.

**Lemma 1.** *For any contextual problem with reward  $\mu$  and context distribution  $\rho$ , let  $\phi$  be a realizable linear representation. The relationship between the diversity conditions in Tab. 2 is summarized in Fig. 3, where each inclusion is in a strict sense and each intersection is non-empty.*

This lemma reveals non-trivial connections between the diversity conditions, better understood through the examples provided in the proof (see App. B.1). BBK is indeed

<sup>4</sup>Whether this is enough for the optimality of the greedy algorithm in the shared-parameter setting is an interesting problem, but it is beyond the scope of this paper.

Name	Definition	Application
Non-redundant	$\lambda_{\min} \left( 1/K \sum_{a \in [K]} \mathbb{E}_{x \sim \rho} [\phi(x, a) \phi(x, a)^\top] \right) > 0$	
CMB	$\forall a, \lambda_{\min} \left( \mathbb{E}_{x \sim \rho} [\phi(x, a) \phi(x, a)^\top] \right) > 0$	Model selection
BBK	$\forall a, \mathbf{u} \in \mathbb{R}^d, \lambda_{\min} \left( \mathbb{E}_x [\phi(x, a) \phi(x, a)^\top \mathbb{1} \{ \phi(x, a)^\top \mathbf{u} \geq 0 \}] \right) > 0$	Logarithmic regret for greedy
HLS	$\lambda_{\min} \left( \mathbb{E}_{x \sim \rho} [\phi^*(x) \phi^*(x)^\top] \right) > 0$	Constant regret for LINUCB
WYS	$\forall a, \lambda_{\min} \left( \mathbb{E}_{x \sim \rho} [\phi(x, a) \phi(x, a)^\top \mathbb{1} \{ x \in \mathcal{X}^*(a) \}] \right) > 0$	Constant regret for LINUCB

Figure 2. Diversity conditions proposed in the literature adapted to the shared-parameter setting. The names refer to the authors who first introduced similar conditions.




Figure 3. Categorization of diversity conditions.

stronger than CMB, and thus it is sufficient for the model selection results by Chatterji et al. (2020). By superficially examining their definitions, CMB may appear stronger than HLS, but the two properties are actually non-comparable, as there are representations that satisfy one condition but not the other. The implications of Fig. 3 on constant-regret guarantees are particularly relevant for our purposes. There are representations that satisfy BBK or CMB and are neither HLS nor WYS and thus may not enable constant regret for LINUCB. We notice that WYS is a stronger condition than HLS. Although WYS may be necessary for LINUCB to achieve constant regret in the disjoint-parameter case, HLS is sufficient for the shared-parameter case we consider in this paper. For this reason, in the following section we adopt HLS to define *good* representations for LINUCB and provide a more complete characterization.

## 4. Good Representations for Constant Regret

The HLS condition was introduced by Hao et al. (2020), who provided a first analysis of its properties. In this section, we complement those results by providing a complete proof of a constant regret bound, a proof of the fact that HLS is actually necessary for constant regret, and a novel characterization of the existence of HLS representations. In the following we define  $\lambda_{\phi, \text{HLS}} := \lambda_{\min} \left( \mathbb{E}_{x \sim \rho} [\phi^*(x) \phi^*(x)^\top] \right)$ , which is strictly positive for HLS representations.

### 4.1. Constant Regret Bound

We begin by deriving a constant problem-dependent regret bound for LINUCB under the HLS condition.

**Lemma 2.** *Consider a contextual bandit problem with realizable linear representation  $\phi$  satisfying the HLS condition (see Tab. 2). Assume  $\Delta > 0$ ,  $\max_{x,a} \|\phi(x, a)\|_2 \leq L$  and  $\|\theta^*\|_2 \leq S$ . Then, with probability at least  $1 - 2\delta$ , the regret of OFUL after  $n \geq 1$  steps is at most*

$$R_n \leq \frac{32\lambda\Delta_{\max}^2 S_\phi^2 \sigma^2}{\Delta} \left( 2 \ln \left( \frac{1}{\delta} \right) + d_\phi \ln \left( 1 + \frac{\tau_\phi L_\phi^2}{\lambda d_\phi} \right) \right)^2,$$

where  $\Delta_{\max} = \max_{x,a} \Delta(x, a)$  is the maximum gap and

$$\tau_\phi \leq \max \left\{ \frac{384^2 d_\phi^2 L_\phi^2 S_\phi^2 \sigma^2 \lambda}{\lambda_{\phi, \text{HLS}} \Delta^2} \ln^2 \left( \frac{64 d_\phi^2 L_\phi^3 \sigma S_\phi \sqrt{\lambda}}{\sqrt{\lambda_{\phi, \text{HLS}} \Delta \delta}} \right), \frac{768 L_\phi^4}{\lambda_{\phi, \text{HLS}}^2} \ln \left( \frac{512 d_\phi L_\phi^4}{\delta \lambda_{\phi, \text{HLS}}^2} \right) \right\}.$$

We first notice that  $\tau_\phi$  is independent from the horizon  $n$ , thus making the previous bound a constant only depending on the problem formulation (i.e., gap  $\Delta$ , norms  $L_\phi$  and  $S_\phi$ ) and the value  $\lambda_{\phi, \text{HLS}}$  which measures “how much” the representation  $\phi$  satisfies the HLS condition. Furthermore, one can always take the minimum between the constant regret in Lem. 2 and any other valid regret bound for OFUL (e.g.,  $O(\log(n)/\Delta)$ ), which may be tighter for small values of  $n$ . While Lem. 2 provides high-probability guarantees, we can easily derive a constant expected-regret bound by running LINUCB with a decreasing schedule for  $\delta$  (e.g.,  $\delta_t \propto 1/t^3$ ) and with a slightly different proof (see App. C and the proof sketch below).

**Proof sketch (full proof in App. C).** Following Hao et al. (2020), the idea is to show that the instantaneous regret  $r_{t+1} = \langle \theta^*, \phi^*(x_{t+1}) - \phi(x_{t+1}, a_{t+1}) \rangle$  is zero for sufficiently large (but constant) time  $t$ . By using the standard regret analysis, we have

$$r_{t+1} \leq 2\beta_{t+1}(\delta) \|\phi(x_{t+1}, a_{t+1})\|_{V_{t+1}^{-1}} \leq \frac{2L\beta_{t+1}(\delta)}{\sqrt{\lambda_{\min}(V_{t+1})}}.$$

Given the minimum-gap assumption, a sufficient condition for  $r_{t+1} = 0$  is that the previous upper bound is smaller than  $\Delta$ , which gives  $\lambda_{\min}(V_{t+1}) > 4L^2\beta_{t+1}^2(\delta)/\Delta^2$ . Since  $\Delta > 0$ , the problem-dependent regret bound in Prop. 1



holds, and the number of pulls to suboptimal arms up to time  $t$  is bounded by  $g_t(\delta) = O((d \ln(t/\delta)/\Delta)^2)$ . Hence, the optimal arms are pulled linearly often and, by leveraging the HLS assumption, we are able to show that the minimum eigenvalue of the design matrix grows linearly in time as

$$\lambda_{\min}(V_{t+1}) \geq \lambda + t\lambda_{\text{HLS}} - 8L^2 \sqrt{t \ln\left(\frac{2dt}{\delta}\right)} - L^2 g_t(\delta).$$

By relating the last two equations, we obtain an inequality of the form  $t\lambda_{\text{HLS}} - o(t) > o(t)$ . If we define  $\tau < \infty$  as the smallest (deterministic) time such that this inequality holds, we have that after  $\tau$  the immediate regret is zero, thus concluding the proof. Note that, if we wanted to bound the expected regret, we could set  $\delta_t \propto 1/t^3$  and the above inequality would still be of the same form (although the resulting  $\tau$  would be slightly different).

**Comparison with existing bounds.** Hao et al. (2020, Theorem 3.9) prove that LINUCB with HLS representations achieves  $\limsup_{n \rightarrow \infty} R_n < \infty$ , without characterizing the time at which the regret vanishes. Instead, our Lem. 2 provides an explicit problem-dependent constant regret bound. Wu et al. (2020, Theorem 2) consider the disjoint-parameter setting and rely on the WYS condition. While they indeed prove a constant regret result, their bound depends on the the minimum probability of observing a context (or, in the continuous case, a properly defined meta-context). This reflects the general tendency, in previous works, to frame diversity conditions simply as a property of the context distribution  $\rho$ . On the other hand, our characterization of  $\tau$  in terms of  $\lambda_{\phi, \text{HLS}}$  (Lem. 2) allows relating the regret to the “goodness” of the representation  $\phi$  for the problem at hand.

## 4.2. Removing the Minimum-Gap Assumption

Constant-regret bounds for LINUCB rely on a minimum-gap assumption ( $\Delta > 0$ ). In this section we show that LINUCB can still benefit from HLS representations when  $\Delta = 0$ , but a margin condition holds (e.g., Rigollet & Zeevi, 2010; Reeve et al., 2018). Intuitively, we require that the probability of observing a context  $x$  decays proportionally to its minimum gap  $\Delta(x) = \min_a \Delta(x, a)$ .

**Assumption 1** (Margin condition). *There exists  $C, \alpha > 2$  such that for all  $\epsilon > 0$ :  $\rho(\{x \in \mathcal{X} : \Delta(x) \leq \epsilon\}) \leq C\epsilon^\alpha$ .*

The following theorem provides a problem-dependent regret bound for LINUCB under this margin assumption.

**Theorem 1.** *Consider a linear contextual bandit problem satisfying the margin condition (Asm. 1). Assume  $\max_{x,a} \|\phi(x, a)\|_2 \leq L_\phi$  and  $\|\theta_\phi^*\|_2 \leq S_\phi$ . Then, given a representation  $\phi$ , with probability at least  $1 - 3\delta$ , the regret of OFUL after  $n \geq 1$  steps is at most*

$$R_n \leq O\left(\left(\lambda(\Delta_{\max} S_\phi \sigma d_\phi)^2 n^{1/\alpha} + \sqrt{Cd_\phi}\right) \ln^2(L_\phi n/\delta)\right).$$

When  $\phi$  is HLS ( $\lambda_{\phi, \text{HLS}} > 0$ ), let  $\tau_\phi \propto (\lambda_{\phi, \text{HLS}})^{\frac{\alpha}{2-\alpha}}$ , then

$$R_n \leq O\left(\Delta_{\max} \tau_\phi + \sqrt{Cd_\phi} \ln^2(L_\phi n/\delta)\right).$$

We first notice that in general, LINUCB suffers  $\tilde{O}(n^{1/\alpha})$  regret, which can be significantly larger than in the minimum-gap case. On the other hand, with HLS representations, LINUCB achieves logarithmic regret, regardless of the value of  $\alpha$ . The intuition is that, when the HLS condition holds, the algorithm collects sufficient information about  $\theta_\phi^*$  by pulling the optimal arms in rounds with large minimum gap, which occur with high probability by the margin condition. This yields at most constant regret in such rounds (first term above), while it can be shown that the regret in steps when the minimum gap is very small is at most logarithmic (second term above).

## 4.3. Further Analysis of the HLS Condition

While Lem. 2 shows that HLS is sufficient for achieving constant regret, the following proposition shows that it is also necessary. While this property was first mentioned by Hao et al. (2020) as a remark in a footnote, we provide a formal proof in App. C.5.

**Proposition 2.** *For any contextual problem with finite contexts, full-support context distribution, and given a non-redundant realizable representation  $\phi$ , LINUCB achieves sub-logarithmic regret if and only if  $\phi$  satisfies the HLS condition.*

As already observed in Section 4, the HLS condition can be equivalently expressed as:<sup>5</sup>

$$\text{span}\{\phi^*(x) \mid x \in \mathcal{X}\} = \mathbb{R}^d,$$

i.e., optimal features must span the whole  $d$ -dimensional Euclidean space, where  $d$  is the dimension of  $\phi$ . If we admit redundant representations, a weaker condition is sufficient to achieve constant regret:

$$\text{span}\{\phi^*(x) \mid x \in \mathcal{X}\} = \text{span}\{\phi(x, a) \mid x \in \mathcal{X}, a \in \mathcal{A}\},$$

i.e., optimal features must span the whole feature space, which may be a subspace of  $\mathbb{R}^d$  in general. We prove that this *weak HLS* condition is sufficient for LINUCB to achieve constant regret as Corollary 1 in App. E.2. This also shows that constant-regret guarantees are preserved by adding redundant features to an HLS representation.

Finally, we derive the following important existence result.

**Lemma 3.** *For any contextual bandit problem with optimal reward <sup>6</sup>  $\mu^*(x) \neq 0$  for all  $x \in \mathcal{X}$ , that has either i) a*

<sup>5</sup>That is assuming the context distribution is full-support. Otherwise, it is enough to replace  $\mathcal{X}$  with the support of the context distribution  $\text{supp}(\rho)$ .

<sup>6</sup>This condition is technical and it can be easily relaxed.

finite context set with at least  $d$  contexts with nonzero probability, or ii) a Borel context space and a non-degenerate context distribution<sup>7</sup>, for any dimension  $d \geq 1$ , there exists an infinite number of  $d$ -dimensional realizable HLS representations.

This result crucially shows that the HLS condition is “robust”, since in any contextual problem, it is possible to construct an infinite number of representations satisfying the HLS condition. In App. B.2, we indeed provide an oracle procedure for constructing an HLS representation. This result also supports the starting point of next section, where we assume that a learner is provided with a set of representations that may contain at least a “good” representation, i.e., an HLS representation.

## 5. Representation Selection

In this section, we study the problem of *representation selection* in linear bandits. We consider a linear contextual problem with reward  $\mu$  and context distribution  $\rho$ . Given a set of  $M$  realizable linear representations  $\{\phi_i : \mathcal{X} \times [K] \rightarrow \mathbb{R}^{d_i}\}$ , the objective is to design a learning algorithm able to perform as well as the best representation, and thus achieve constant regret when a “good” representation is available. As usual, we assume  $\theta_i^* \in \mathbb{R}^{d_i}$  is unknown, but the algorithm is provided with a bound on the parameter and feature norms of the different representations.

### 5.1. The LEADER Algorithm

We introduce LEADER (*Linear rEpresentation bAnDit mixER*), see Alg. 1. At each round  $t$ , LEADER builds an estimate  $\theta_{ti}$  of the unknown parameter  $\theta_i^*$  of each representation  $\phi_i$ .<sup>8</sup> These estimates are by nature off-policy, and thus *all* the samples  $(x_l, a_l, y_l)_{l < t}$  can be used to solve *all* ridge regression problems. For each  $\phi_i$ , define  $V_{ti} = \lambda I_{d_i} + \sum_{l=1}^{t-1} \phi_i(x_l, a_l) \phi_i(x_l, a_l)^\top$ ,  $\theta_{ti}$  and  $\mathcal{C}_{ti}(\delta/M)$  as in Sec. 2. Since all the representations are realizable, we have that  $\mathbb{P}(\forall i \in [M], \theta_i^* \in \mathcal{C}_{ti}(\delta/M)) \geq 1 - \delta$ . As a consequence, for each representation  $\phi_i$  we can build an upper-confidence bound to the reward such that,  $\forall x \in \mathcal{X}, a \in \mathcal{A}$ , with high probability

$$\mu(x, a) \leq \max_{\theta \in \mathcal{C}_{ti}(\delta/M)} \langle \phi_i(x, a), \theta \rangle := U_{ti}(x, a). \quad (1)$$

Given this, LEADER uses the tightest available upper-confidence bound to evaluate each action and then it selects the one with the largest value, i.e.,

$$a_t \in \operatorname{argmax}_{a \in [K]} \min_{i \in [M]} \{U_{ti}(x_t, a)\}. \quad (2)$$

<sup>7</sup>For instance, if  $\mathcal{X} = \mathbb{R}^m$  and the context distribution must have positive variance in all directions.

<sup>8</sup>We use the subscript  $i \in [M]$  instead of  $\phi_i$  to denote quantities related to representation  $\phi_i$ .

Let  $i_t = \operatorname{argmin}_{i \in [M]} \{U_{ti}(x_t, a_t)\}$  be the representation associated to the pulled arm  $a_t$ . Interestingly, despite  $a_t$  being optimistic, in general it may not correspond to the optimistic action of representation  $\phi_{i_t}$ , i.e.,  $a_t \notin \operatorname{argmax}_a \{U_{t, i_t}(x_t, a)\}$ . If a representation provides an estimate that is good along the direction associated to a context-action pair, but possibly very uncertain on other actions, LEADER is able to leverage this key feature to reduce the overall uncertainty and achieve a tighter optimism. Space and time complexity of LEADER scales linearly in the number of representations, although the updates for each representation could be carried out in parallel.

**Regret bound.** For ease of presentation, we assume a non-zero minimum gap ( $\Delta > 0$ ). The analysis can be generalized to  $\Delta = 0$  as done in Sec. 4.2. Thm. 2 establishes the regret guarantee of LEADER (Alg. 1).

**Theorem 2.** *Consider a contextual bandit problem with reward  $\mu$ , context distribution  $\rho$  and  $\Delta > 0$ . Let  $(\phi_i)$  be a set of  $M$  linearly realizable representations such that  $\max_{x, a} \|\phi_i(x, a)\|_2 \leq L_i$  and  $\|\theta_i^*\| \leq S_i$ . Then, for any  $n \geq 1$ , with probability  $1 - 2\delta$ , LEADER suffers a regret*

$$R_n \leq \min_{i \in [M]} \left\{ \frac{32\lambda\Delta_{\max}^2 S_i^2 \sigma^2}{\Delta} \times \left( 2 \ln \left( \frac{M}{\delta} \right) + d_i \ln \left( 1 + \frac{\min\{\tau_i, n\} L_i^2}{\lambda d_i} \right) \right)^2 \right\}$$

where  $\tau_i \propto (\lambda_{i, \text{HLS}} \Delta)^{-2}$  if  $\phi_i$  is HLS and  $\tau_i = +\infty$  otherwise.

This shows that the problem-dependent regret bound of LEADER is not worse than the one of the best representation (see Prop. 1), up to a  $\ln M$  factor. This means that the cost of representation selection is almost negligible. Furthermore, Thm. 2 shows that LEADER not only achieves a constant regret bound when an HLS representation is available, but this bound scales as the one of the *best* HLS representation. In fact, notice that the “quality” of an HLS representation does not depend only on known quantities such as  $d_i, L_i, S_i$ , but crucially on HLS eigenvalue  $\lambda_{i, \text{HLS}}$ , which is usually not known in advance, as it depends on the features of the optimal arms.

### 5.2. Combining Representations

In the previous section, we have shown that LEADER can perform as well as the best representation in the set. However, by inspecting the action selection rule (Eq. 2), we notice that, to evaluate the reward of an action in the current context, LEADER selects the representation with the smallest uncertainty, thus potentially using different representations for different context-action pairs. This leads to the question: *can LEADER do better than the best representation in the set?*

**Algorithm 1** LEADER Algorithm

---

**Input:** representations  $(\phi_i)_{i \in [M]}$  with values  $(L_i, S_i)_{i \in [M]}$ , regularization factor  $\lambda \geq 1$ , confidence level  $\delta \in (0, 1)$ .  
 Initialize  $V_{1i} = \lambda I_{d_i}$ ,  $\theta_{1i} = 0_{d_i}$  for each  $i \in [M]$   
**for**  $t = 1, \dots$  **do**  
   Observe context  $x_t$   
   Pull action  $a_t \in \operatorname{argmax}_{a \in [K]} \min_{i \in [M]} \{U_{ti}(x_t, a)\}$   
   Observe reward  $r_t$  and, for each  $i \in [M]$ , set  
    $V_{t+1,i} = V_{ti} + \phi_i(x_t, a_t) \phi_i(x_t, a_t)^\top$  and  
    $\theta_{t+1,i} = V_{t+1,i}^{-1} \sum_{s=1}^t \phi_i(x_s, a_s) r_s$   
**end for**

---

We show that, in certain cases, LEADER is able to combine representations and achieve constant regret when none of the individual representations would. The intuition is that a subset of “locally good” representations can be combined to recover a condition similar to HLS. This property is formally stated in the following definition.

**Definition 1** (Mixing HLS). *Consider a linear contextual problem with reward  $\mu$  and context distribution  $\rho$ , and a set of  $M$  realizable linear representations  $\phi_1, \dots, \phi_M$ . Define  $M_i = \mathbb{E}_{x \sim \rho} [\phi_i^*(x) \phi_i^*(x)^\top]$  and let  $Z_i = \{(x, a) \in \mathcal{X} \times \mathcal{A} \mid \phi_i(x, a) \in \operatorname{Im}(M_i)\}$  be the set of context-action pairs whose features belong to the column space of  $M_i$ , i.e., that lie in the span of optimal features. We say that the set  $(\phi_i)$  satisfies the mixed-HLS condition if  $\mathcal{X} \times \mathcal{A} \subseteq \bigcup_{i=1}^M Z_i$ .*

Let  $\lambda_i^+ = \lambda_{\min}^+(M_i)$  be the minimum *nonzero* eigenvalue of  $M_i$ . Intuitively, the previous condition relies on the observation that every representation satisfies a “restricted” HLS condition on the context-action pairs  $(x, a)$  whose features  $\phi_i(x, a)$  are spanned by optimal features  $\phi^*(x)$ . In this case, the characterizing eigenvalue is  $\lambda_i^+$ , instead of the smallest eigenvalue  $\lambda_{i,\text{HLS}}$  (which may be zero). If every context-action pair is in the restriction  $Z_i$  of some representation, we have the mixed-HLS property. In particular, if representation  $i$  is HLS,  $\lambda_i^+ = \lambda_{i,\text{HLS}}$  and  $Z_i = \mathcal{S} \times \mathcal{A}$ . So, HLS is a special case of mixed-HLS. In App. E.2, we provide simple examples of sets of representations satisfying Def. 1. Note that, strictly speaking, there is not a single “mixed representation” solving the whole problem. Even defining one would be problematic since each representation may have a different parameter and even a different dimension. Instead, each representation “specializes” on a different portion of the context-action space. If together they cover the whole space, the benefits of HLS are recovered, as illustrated in the following theorem.

**Theorem 3.** *Consider a stochastic bandit problem with reward  $\mu$ , context distribution  $\rho$  and  $\Delta > 0$ . Let  $(\phi_i)$  be a set of  $M$  realizable linear representations satisfying the mixed-HLS property in Def. 1. Then, with probability at least  $1 - 2\delta$ , there exists a time  $\tau < \infty$  independent from  $n$  such that, for any  $n \geq 1$ , the pseudo-regret of LEADER is*

bounded as

$$R_n \leq \min_{i \in [M]} \left\{ \frac{32\lambda\Delta_{\max}^2 S_i^2 \sigma^2}{\Delta} \times \left( 2 \ln \left( \frac{M}{\delta} \right) + d_i \ln \left( 1 + \frac{\tau L_i^2}{\lambda d_i} \right) \right)^2 \right\}.$$

First, note that we are still scaling with the characteristics of the best representation in the set (i.e.,  $d_i$ ,  $L_i$  and  $S_i$ ). However, the time  $\tau$  to constant regret is a global value rather than being different for each representation. This highlights that mixed-HLS is a global property of the set of representations rather than being individual as before. In particular, whenever no representation is (globally) HLS (i.e.,  $\lambda_{i,\text{HLS}} = 0$  for all  $\phi_i$ ), we can show that in the worst case  $\tau$  scales as  $(\min_i \lambda_i^+)^{-2}$ . In practice, we may expect LEADER to even behave better than that since **i**) not all the representations may contribute actively to the mixed-HLS condition; and **ii**) multiple representations may cover the same region of the context-action space. In the latter case, since LEADER leverages all the representations at once, its regret would rather scale with the largest minimum nonzero eigenvalue  $\lambda_i^+$  among all the representations covering such region. We refer to App. E.2 for a more complete discussion.

### 5.3. Discussion

Most of the model selection algorithms reviewed in the introduction could be readily applied to select the best representation for LINUCB. However, the generality of their objective comes with several shortcomings when instantiated in our specific problem (see App. A for a more detailed comparison). First, model selection methods achieve the performance of the best algorithm, up to a polynomial dependence on the number  $M$  of models. This already makes them a weaker choice compared to LEADER, which, by leveraging the specific structure of the problem, suffers only a logarithmic dependence on  $M$ . Second, model selection algorithms are often studied in a worst-case analysis, which reveals a high cost for adaptation. For instance, corollary algorithms (Agarwal et al., 2017; Pacchiano et al., 2020b) pay an extra  $\sqrt{n}$  regret, which would make them unsuitable to target the constant regret of good representations. Similar costs are common to other approaches (Abbasi-Yadkori et al., 2020; Pacchiano et al., 2020a). It is unclear whether a problem-dependent analysis can be carried out and whether this could shave off such dependence. Third, these algorithms are generally designed to adapt to a specific best base algorithm. At the best of our knowledge, there is no evidence that model selection methods could combine algorithms to achieve better performance than the best candidate, a behavior that we proved for LEADER in our setting.

On the other hand, model selection algorithms effec-

tively deal with non-realizable representations in certain cases (e.g., Foster et al., 2020; Abbasi-Yadkori et al., 2020; Pacchiano et al., 2020a), while LEADER is limited to the realizable case. While a complete study of the model misspecification case is beyond the scope of this paper, in App. F, we discuss how a variation of the approach presented in (Agarwal et al., 2012b) could be paired to LEADER to discard misspecified representations and possibly recover the properties of “good” representations.

## 6. Experiments

In this section, we report experimental results on two synthetic and one dataset-based problems. For each problem, we evaluate the behavior of LEADER with LINUCB and model selection algorithms: EXP4.IX (Neu, 2015), CORRAL and EXP3.P in the stochastic version by Pacchiano et al. (2020b) and Regret Balancing with and without elimination (REGBALELIM and REGBAL) (Abbasi-Yadkori et al., 2020; Pacchiano et al., 2020a). See App. G for a detailed discussion and additional experiments. All results are averaged over 20 independent runs, with shaded areas corresponding to 2 standard deviations. We always set the parameters to  $\lambda = 1$ ,  $\delta = 0.01$ , and  $\sigma = 0.3$ . All the representations we consider are normalized to have  $\|\theta_i^*\| = 1$ .

**Synthetic Problems.** We define a randomly-generated contextual bandit problem, for which we construct sets of realizable linear representations with different properties (see App. G.1 for details). The purpose of these experiments is twofold: to show the different behavior of LINUCB with different representations, and to evaluate the ability of LEADER of selecting and mixing representations.

*Varying dimension.* We construct six representations of varying dimension from 2 up to 6. Of the two representations of dimension  $d = 6$ , one is HLS. Fig. 4(left) shows that in this case, LINUCB with the HLS representation outperforms any non-HLS representation, even if they have smaller dimension. This property is inherited by LEADER, which performs better than LINUCB with non-HLS representations even of much smaller dimension 2.

*Mixing representations.* We construct six representations of the same dimension  $d = 6$ , none of which is HLS. However, they are constructed so that together they satisfy the weaker mixed-HLS assumption (Def. 1). Fig. 4(middle left) shows that, as predicted by Thm. 3, LEADER leverages different representations in different context-action regions and it thus performs significantly better than any LINUCB using non-HLS representations. The superiority of LEADER w.r.t. the model-selection baselines is evident in this case (Fig. 4(middle right)), since only LEADER is able to mix representations, whereas model-selection algorithms target the best in a set of “bad” representations. Additional experiments in App. G confirm that LEADER consistently

outperforms all model-selection algorithms.

**Jester Dataset.** In the last experiment, we extract multiple linear representations from the Jester dataset (Goldberg et al., 2001), which consists of joke ratings in a continuous range from  $-10$  to  $10$  for a total of 100 jokes and 73421 users. For a subset of 40 jokes and 19181 users rating all these 40 jokes, we build a linear contextual problem as follows. First, we fit a  $32 \times 32$  neural network to predict the ratings from features extracted via a low-rank factorization of the full matrix. Then, we take the last layer of the network as our “ground truth” linear model and fit multiple smaller networks to clone its predictions, while making sure that the resulting misspecification is small. We thus obtain 7 representations with different dimensions among which, interestingly, we find that 6 are HLS. Figure 4(right) reports the comparison between LEADER using all representations and LINUCB with each single representation on a log-scale. Notably, the ability of LEADER to mix representations makes it perform better than the best candidate, while transitioning to constant regret much sooner. Finally, the fact that HLS representations arise so “naturally” raises the question of whether this is a more general pattern in context-action features learned from data.

**Last.fm dataset.** In App. F we study a variant of LEADER that is able to handle misspecified representations, and we test it on the Last.fm music-recommendation dataset (Cantador et al., 2011). See App. F.4 for details.

## 7. Conclusion

We provided a complete characterization of “good” realizable representations for LINUCB, ranging from existence to a sufficient and necessary condition to achieve problem-dependent constant regret. We introduced LEADER, a novel algorithm that, given a set of realizable linear representations, is able to adapt to the best one and even leverage their combination to achieve constant regret under the milder mixed-HLS condition. While we have focused on LINUCB, other algorithms (e.g., LinTS (Abeille & Lazaric, 2017)) as well as other settings (e.g., low-rank RL (Jin et al., 2020)) may also benefit from HLS-like assumptions. We have mentioned an approach for eliminating misspecified representations, but a non-trivial trade-off may exist between the level of misspecification and the goodness of the representation. A slightly imprecise but very informative representation may be preferable to most bad realizable ones. Finally, we believe that moving from selection to representation learning –e.g., provided a class of features such as a neural network– is an important direction both from a theoretical and practical perspective.






Figure 4. Regret of LEADER and model-selection baselines on different linear contextual bandit problems. (left) Synthetic problem with varying dimensions. (middle left) Representation mixing. (middle right) Comparison to model selection baselines. (right) Jester dataset.

## References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. In *NIPS*, pp. 2312–2320, 2011.
- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *AISTATS*, volume 22 of *JMLR Proceedings*, pp. 1–9. JMLR.org, 2012.
- Abbasi-Yadkori, Y., Pacchiano, A., and Phan, M. Regret balancing for bandit and rl model selection, 2020.
- Abeille, M. and Lazaric, A. Linear thompson sampling revisited. In *AISTATS*, volume 54 of *Proceedings of Machine Learning Research*, pp. 176–184. PMLR, 2017.
- Agarwal, A., Dudík, M., Kale, S., Langford, J., and Schapire, R. Contextual bandit learning with predictable rewards. In *Artificial Intelligence and Statistics*, pp. 19–26. PMLR, 2012a.
- Agarwal, A., Dudík, M., Kale, S., Langford, J., and Schapire, R. E. Contextual bandit learning with predictable rewards. In *AISTATS*, volume 22 of *JMLR Proceedings*, pp. 19–26. JMLR.org, 2012b.
- Agarwal, A., Luo, H., Neyshabur, B., and Schapire, R. E. Corraling a band of bandit algorithms. In *COLT*, volume 65 of *Proceedings of Machine Learning Research*, pp. 12–38. PMLR, 2017.
- Arora, R., Marinov, T. V., and Mohri, M. Corraling stochastic bandit algorithms. *CoRR*, abs/2006.09255, 2020.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- Bastani, H., Bayati, M., and Khosravi, K. Mostly exploration-free algorithms for contextual bandits. *Management Science*, 2020.
- Bibaut, A. F., Chambaz, A., and van der Laan, M. J. Rate-adaptive model selection over a collection of black-box contextual bandit algorithms. *CoRR*, abs/2006.03632, 2020.
- Bouneffouf, D. and Rish, I. A survey on practical applications of multi-armed and contextual bandits, 2019.
- Cantador, I., Brusilovsky, P., and Kuflik, T. 2nd workshop on information heterogeneity and fusion in recommender systems (hetrec 2011). In *Proceedings of the 5th ACM conference on Recommender systems, RecSys 2011*, New York, NY, USA, 2011. ACM. URL <http://ir.ii.uam.es/hetrec2011/index.html>. Last.fm website, <http://www.lastfm.com>.
- Chatterji, N. S., Muthukumar, V., and Bartlett, P. L. OSOM: A simultaneously optimal algorithm for multi-armed and linear contextual bandits. In *AISTATS*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1844–1854. PMLR, 2020.
- Chatzigeorgiou, I. Bounds on the lambert function and their application to the outage analysis of user cooperation. *CoRR*, abs/1601.04895, 2016.
- Chen, F. A note on matrix versions of kantorovich-type inequality. *JOURNAL OF MATHEMATICAL INEQUALITIES*, 7(2):283–288, 2013.
- Chu, W., Li, L., Reyzin, L., and Schapire, R. E. Contextual bandits with linear payoff functions. In *AISTATS*, volume 15 of *JMLR Proceedings*, pp. 208–214. JMLR.org, 2011.
- Degenne, R., Ménard, P., Shang, X., and Valko, M. Gamification of pure exploration for linear bandits. In *International Conference on Machine Learning*, pp. 2432–2442. PMLR, 2020.
- Eaton, M. L. and Perlman, M. D. The non-singularity of generalized sample covariance matrices. *The Annals of Statistics*, pp. 710–717, 1973.
- Foster, D. J., Krishnamurthy, A., and Luo, H. Model selection for contextual bandits. In *NeurIPS*, pp. 14714–14725, 2019.

- Foster, D. J., Gentile, C., Mohri, M., and Zimmert, J. Adapting to misspecification in contextual bandits. In *NeurIPS*, 2020.
- Ghosh, A., Sankararaman, A., and Ramchandran, K. Problem-complexity adaptive model selection for stochastic linear bandits. *CoRR*, abs/2006.02612, 2020.
- Goldberg, K. Y., Roeder, T., Gupta, D., and Perkins, C. Eigentaste: A constant time collaborative filtering algorithm. *Inf. Retr.*, 4(2):133–151, 2001.
- Hao, B., Lattimore, T., and Szepesvári, C. Adaptive exploration in linear contextual bandit. In *AISTATS*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3536–3545. PMLR, 2020.
- Hoorfar, A. and Hassani, M. Inequalities on the lambert w function and hyperpower function. *J. Inequal. Pure and Appl. Math.*, 9(2):5–9, 2008.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *COLT*, volume 125 of *Proceedings of Machine Learning Research*, pp. 2137–2143. PMLR, 2020.
- Lattimore, T. and Szepesvari, C. The end of optimism? an asymptotic analysis of finite-armed linear bandits. In *Artificial Intelligence and Statistics*, pp. 728–737. PMLR, 2017.
- Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2020.
- Lattimore, T., Szepesvari, C., and Weisz, G. Learning with good feature representations in bandits and in rl with a generative model. In *International Conference on Machine Learning*, pp. 5662–5670. PMLR, 2020.
- Lee, J. N., Pacchiano, A., Muthukumar, V., Kong, W., and Brunskill, E. Online model selection for reinforcement learning with function approximation, 2020.
- Maillard, O. and Munos, R. Adaptive bandits: Towards the best history-dependent strategy. In *AISTATS*, volume 15 of *JMLR Proceedings*, pp. 570–578. JMLR.org, 2011.
- Neu, G. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. In *NeurIPS*, pp. 3168–3176, 2015.
- Pacchiano, A., Dann, C., Gentile, C., and Bartlett, P. Regret bound balancing and elimination for model selection in bandits and rl, 2020a.
- Pacchiano, A., Phan, M., Abbasi-Yadkori, Y., Rao, A., Zimmert, J., Lattimore, T., and Szepesvári, C. Model selection in contextual stochastic bandit problems. In *NeurIPS*, 2020b.
- Reeve, H. W. J., Mellor, J., and Brown, G. The k-nearest neighbour UCB algorithm for multi-armed bandits with covariates. In *ALT*, volume 83 of *Proceedings of Machine Learning Research*, pp. 725–752. PMLR, 2018.
- Rigollet, P. and Zeevi, A. Nonparametric bandits with covariates. *arXiv preprint arXiv:1003.1630*, 2010.
- Tirinzi, A., Pirota, M., Restelli, M., and Lazaric, A. An asymptotically optimal primal-dual incremental algorithm for contextual linear bandits. *Advances in Neural Information Processing Systems*, 33, 2020.
- Tropp, J. A. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, 12(4):389–434, 2012.
- Wu, W., Yang, J., and Shen, C. Stochastic linear contextual bandits with diverse contexts. In *AISTATS*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2392–2401. PMLR, 2020.

## A. Comparison to Related Work

In this section, we provide a more detailed review of the literature of model-selection in contextual bandits. Note that this literature has mainly focused on analyzing the minimax regime. Studying whether it is possible and/or how to leverage a problem-dependent analysis is outside the scope of this work. We thus present a review of the theoretical results of these algorithms, once applied to the representation selection problem considered in this paper.

In general, a model-selection algorithm has access to a set of  $M$  base contextual-bandit algorithms (in the following, simply *bases*), and, at each step, it selects a base and plays the corresponding action. Our bases are instances of LINUCB with different (realizable) representations.

One of the representative algorithms in model selection is CORRAL (Agarwal et al., 2017). However, the stability conditions required by CORRAL are not satisfied by LINUCB. For this reason, we consider the stochastic version in (Pacchiano et al., 2020b), which can be applied to any base algorithm after applying a *smoothing wrapper*. This approach requires the knowledge of an upper bound to the regret of the best base algorithm. For simplicity, we consider the  $\tilde{O}(\sqrt{n})$  worst-case bound (e.g., Abbasi-Yadkori et al., 2011, Thm. 3). Let  $c(\delta)$  be the leading constant from that bound. Neglecting logarithmic terms and setting the initial learning rate of CORRAL to  $\eta = \sqrt{M/n}/c(\delta)$ , according to Theorem 5.3. from Pacchiano et al. (2020b):

$$R_n \leq \tilde{O} \left( (1 + c(\delta))\sqrt{Mn} + \frac{\sqrt{Mn}}{c(\delta)} \right).$$

Pacchiano et al. (2020b) showed a minimax lower bound  $\Omega(T)$  for adversarial masters. Despite being a minimax result, whether the inherent  $\sqrt{T}$  of adversarial masters can be shaved off through a refined analysis is an open question, as well as how to leveraged more adaptive regret bounds. Another major downside of CORRAL is the  $\sqrt{M}$  dependence on the number of representations, compared to the logarithmic dependence of LEADER. Clearly, these downsides are compensated by the more general applicability of these algorithms. See (Arora et al., 2020; Foster et al., 2020) for examples of usage of corraling techniques. Pacchiano et al. (2020b) also propose to use EXP3.P as a master algorithm in place of CORRAL, using the same smoothing wrapper for base algorithms and the same regret oracle. By setting the explicit-exploration parameter of EXP3.P to  $p = n^{-1/3}M^{-2/3}c(\delta)^{2/3}$ , according to Theorem 5.3. from Pacchiano et al. (2020b):  $R_n \leq \tilde{O} \left( \sqrt{Mn} + M^{1/3}n^{2/3}c(\delta)^{2/3} \right)$ . The same considerations made for CORRAL apply to this case.

Recently, techniques based on the idea of regret balancing has been used for model selection (Abbasi-Yadkori et al., 2020; Pacchiano et al., 2020a). Denote by  $\mathcal{H}_{ti} = (x_k, a_k, r_k)_{k \in N_i(t)}$  as the history associated to base algorithm  $i$ , i.e., restricted to rounds where  $i$  was selected. The Regret Balancing algorithm by Abbasi-Yadkori et al. (2020) (REGBAL for short) requires access to a (high-probability, possibly data dependent) upper bound  $u : \mathcal{H}_{ti^*} \mapsto \mathbb{R}$  on the regret of the best base algorithm  $i^* \in [M]$ . The regret of REGBAL is then, with high probability:

$$R_n \leq M \max_{i \in [M]} u(\mathcal{H}_{ni}),$$

that is, bounded like the best base on the worst subjective history. While this bound is very implicit, it shows: *i*) REGBAL pays a linear dependence on the number of base algorithms; *ii*) even using a problem dependent bound, the final result will scale as the worse base algorithm.

The second regret balancing approach is by (Pacchiano et al., 2020a). Their Regret Balancing and Elimination by Pacchiano et al. (2020a) (REGBALELIM for short) requires an upper bound  $u_i : \mathcal{H}_i \mapsto \mathbb{R}$  on the regret of *each* base algorithm (can be different for each algorithm). Their algorithms is able to eliminate the bases for which  $u_i$  is incorrectly specified, and compete with the best *correct* upper bound. Their algorithm improves w.r.t. REGBAL and, by simplifying a lot, may allow to scale as the regret of the best base algorithm. However, *i*) the dependence on  $M$  is still linear and *ii*) whether the idea of regret balancing will allow to break the barrier of logarithmic regret is an open question. As shown in their paper (e.g., Table 1), even if the regret upper bound is  $O(n^\beta)$  with  $\beta < 1/2$ , their bound has a (minimax)  $O(\sqrt{n})$  dependence.

Finally, we would like to mention also the literature about learning with expert advice. In particular, EXP4.IX (Neu, 2015) is an algorithm for bandits with expert advice, itself an improvement over the original EXP4 (Auer et al., 2002). At each step, each expert prescribes an action (more in general, a distribution over actions) based on the whole history. EXP4.IX observes the prescriptions of all the experts and decides which action to play. In our case, the experts are the bases. Similarly to LEADER, EXP4.IX is not strictly a model-selection algorithm since the action it plays may be different from the one

prescribed by any base. In our notation, the (worst-case, high-probability) regret bound from (Neu, 2015, Thm. 2) is:

$$R_n \leq 2\sqrt{2Kn \log M} + \left( \sqrt{\frac{2Kn}{\ln n}} + 1 \right) \ln(2/\delta) + u_n,$$

where  $u_n$  is an upper bound on the regret of LINUCB with the best representation. Even in the case the best representation is HLS (i.e.,  $u_n$  is a constant), the regret bound is still  $\tilde{O}(\sqrt{n})$ . Moreover, EXP4.IX has an explicit dependence on the number of actions  $K$ . An upside is the logarithmic dependence on the number of representations  $M$ .

## B. Analysis of Several Diversity Conditions

### B.1. Categorization

In this Section, we provide a complete proof of Lemma 1. The (counter)examples we provide to prove that the intersections depicted in Figure 3 are nonempty, we believe, can also improve understanding of the diversity conditions.

In comparing representations, we fix a contextual problem with reward  $\mu$  and context distribution  $\rho$ , and a dimension  $d \geq 1$ . We use set notation for conciseness.<sup>9</sup> First, we provide in Table 1 equivalent definitions, in terms of the span of particular sets of features, of the properties defined in Figure 2. Without loss of generality, we assume that the context distribution  $\rho$  is full-support, i.e.,  $\text{supp}(\rho) = \mathcal{X}$ . Otherwise, it is enough to replace  $\mathcal{X}$  with  $\text{supp}(\rho)$  in the definitions.

Name	Definition
Non-redundant	$\text{span} \{ \phi(x, a) \mid x \in \mathcal{X}, a \in \mathcal{A} \} = \mathbb{R}^d$
CMB	$\forall a, \text{span} \{ \phi(x, a) \mid x \in \mathcal{X} \} = \mathbb{R}^d$
BBK	$\forall a, \mathbf{u} \in \mathbb{R}^d, \text{span} \{ \phi(x, a) \mid x \in \mathcal{X}, \phi(x, a)^\top \mathbf{u} \geq 0 \} = \mathbb{R}^d$
HLS	$\text{span} \{ \phi^*(x) \mid x \in \mathcal{X} \} = \mathbb{R}^d$
WYS	$\forall a, \text{span} \{ \phi(x, a) \mid x \in \mathcal{X}^*(a) \} = \mathbb{R}^d$

Table 1. Diversity conditions in terms of the span of particular sets of features.

The examples we provide are always 2-dimensional representations with a corresponding parameter  $\theta^* = [1, 1]$ , and refer to the contextual problem having the resulting reward function and uniform context distribution. Optimal features are underlined for convenience.

**CMB  $\subset$  non-redundant.** For any  $a \in \mathcal{A}$ ,  $\text{span} \{ \phi(x, a) \mid x \in \mathcal{X} \} \subseteq \text{span} \{ \phi(x, a') \mid x \in \mathcal{X}, a' \in \mathcal{A} \}$ . In words, if the features of a single arm span  $\mathbb{R}^d$ , so does the whole set of features. Not all non-redundant representations are CMB, as testified by the following representation:

$$\begin{aligned} \underline{\phi(x_1, a_1)} &= [1, 1] & \phi(x_1, a_2) &= [1/2, 1/2] \\ \phi(x_2, a_1) &= [0, 1] & \underline{\phi(x_2, a_2)} &= [1, 1] \end{aligned}$$

which is non-redundant since all features together span  $\mathbb{R}^2$ , but not CMB since those of arm  $a_2$  alone fail to do so.

**HLS  $\subset$  Non-redundant.** Since  $\phi^*(x) = \phi^*(x, a_x^*)$  for some  $a_x^* \in \mathcal{A}$ ,  $\text{span} \{ \phi^*(x) \mid x \in \mathcal{X} \} \subseteq \text{span} \{ \phi(x, a') \mid x \in \mathcal{X}, a' \in \mathcal{A} \}$ . In words, if optimal features span  $\mathbb{R}^d$ , so does the whole set of features. Not all non-redundant representations are HLS, as testified by the following representation:

$$\begin{aligned} \underline{\phi(x_1, a_1)} &= [1, 1] & \phi(x_1, a_2) &= [1/2, 1/2] \\ \phi(x_2, a_1) &= [0, 1] & \underline{\phi(x_2, a_2)} &= [1, 1] \end{aligned}$$

which is non-redundant since all features together span  $\mathbb{R}^2$ , but not HLS since optimal features alone fail to do so.

<sup>9</sup>For instance, we use HLS to denote the set of  $d$ -dimensional realizable representations that are HLS for the given contextual problem.



**CMB  $\not\subseteq$  HLS.** The following representation is CMB but not HLS:

$$\begin{aligned} \underline{\phi(x_1, a_1)} &= [1, 1] & \phi(x_1, a_2) &= [1, 0] \\ \underline{\phi(x_2, a_1)} &= [0, 1] & \underline{\phi(x_2, a_2)} &= [1, 1] \end{aligned} \quad (3)$$

since the features of each arm taken alone span  $\mathbb{R}^2$ , but optimal features fail to do so.

**HLS  $\not\subseteq$  CMB.** The following representation is HLS but not CMB:

$$\begin{aligned} \underline{\phi(x_1, a_1)} &= [2, 0] & \phi(x_1, a_2) &= [1/2, 1/2] \\ \underline{\phi(x_2, a_1)} &= [0, 2] & \phi(x_2, a_2) &= [1/2, 1/2] \end{aligned}$$

since optimal features span  $\mathbb{R}^2$ , but the features of arm  $a_2$  alone fail to do so.

**WYS  $\subset$  CMB  $\cap$  HLS.** Assume  $\phi$  is WYS. For all  $a$ , since  $\mathcal{X}^*(a) \subseteq \mathcal{X}$ ,  $\text{span}\{\phi(x, a) \mid x \in \mathcal{X}^*(a)\} \subseteq \text{span}\{\phi(x, a) \mid x \in \mathcal{X}\}$ . Hence  $\phi$  is CMB. If  $x \in \mathcal{X}^*(a)$ , then  $\phi(x, a) = \phi^*(x)$ . So, for any  $a$ ,  $\text{span}\{\phi(x, a) \mid x \in \mathcal{X}^*(a)\} \subseteq \text{span}\{\phi^*(x) \mid x \in \mathcal{X}\}$ . Hence  $\phi$  is HLS. The inclusion is strict, as testified by the following representation, which is both HLS and CMB:

$$\begin{aligned} \underline{\phi(x_1, a_1)} &= [2, 0] & \phi(x_1, a_2) &= [1, 0] \\ \underline{\phi(x_2, a_1)} &= [0, 2] & \phi(x_2, a_2) &= [0, 1] \end{aligned}$$

but is not WYS since  $a_1$  is only optimal for  $x_1$ , and  $\phi(x_1, a_1)$  cannot span  $\mathbb{R}^2$  alone.

**CMB  $\cap$  HLS  $\neq \emptyset$ .** See the previous example.

**BBK  $\subset$  CMB.** For all action  $a$  and  $\mathbf{u} \in \mathbb{R}^d$ ,  $\text{span}\{\phi(x, a) \mid x \in \mathcal{X}, \phi(x, a)^\top \mathbf{u} \geq 0\} \subseteq \text{span}\{\phi(x, a) \mid x \in \mathcal{X}\}$ . So if a representation is BBK, it is also CMB. The converse is not true in general: none of the examples we have provided so far are BBK, since all features lie in the first quadrant of  $\mathbb{R}^2$ , hence  $\text{span}\{\phi(x, a) \mid x \in \mathcal{X}, \phi(x, a)^\top \mathbf{u} \geq 0\} = \emptyset$  for  $\mathbf{u} = [-1, -1]$ . In particular, (3) is CMB but not BBK. This, together with HLS  $\not\subseteq$  CMB, also shows that HLS  $\not\subseteq$  BBK.

**BBK  $\not\subseteq$  HLS** The following representation:

$$\begin{aligned} \underline{\phi(x_1, a_1)} &= [2, 0] & \phi(x_1, a_2) &= [0, 1] \\ \phi(x_2, a_1) &= [0, 1] & \underline{\phi(x_2, a_2)} &= [2, 0] \\ \underline{\phi(x_2, a_1)} &= [-1, 0] & \phi(x_2, a_2) &= [0, -2] \\ \phi(x_2, a_1) &= [0, -2] & \underline{\phi(x_2, a_2)} &= [-1, 0] \end{aligned} \quad (4)$$

is BBK but not HLS. To show that it is not HLS, we just notice that optimal features fail to span  $\mathbb{R}^2$ . To show that it is BBK, we can easily check that  $\{\phi(x, a) \mid x \in \mathcal{X}, \phi(x, a)^\top \mathbf{u} \geq 0\}$  spans  $\mathbb{R}^2$  for both actions and  $\mathbf{u} \in \{[1, 1], [-1, 1], [-1, -1], [1, -1]\}$ . Any other vector can be obtained as  $A\mathbf{u}$  from one of these four vectors, where  $A$  is a p.s.d. diagonal matrix. Since all features  $\phi$  are aligned with the axes of  $\mathbb{R}^2$ ,  $\phi^\top \mathbf{u} = \phi_i \times u_i$ , for some index  $i \in \{1, 2\}$ . Similarly,  $\mathbb{R}^d$ ,  $\phi(x, a)^\top A\mathbf{u} = \phi_i \times A_{ii} \times u_i$ . Since  $A_{ii} \geq 0$ , non-negativity of the former guarantees non-negativity of the latter. This example also shows that BBK is not empty in general.

**HLS  $\cap$  BBK  $\not\subseteq$  WYS.** The following representation:

$$\begin{aligned} \underline{\phi(x_1, a_1)} &= [2, 0] & \phi(x_1, a_2) &= [1, 0] \\ \underline{\phi(x_2, a_1)} &= [0, 2] & \phi(x_2, a_2) &= [0, 1] \\ \underline{\phi(x_2, a_1)} &= [-1, 0] & \phi(x_2, a_2) &= [-2, 0] \\ \underline{\phi(x_2, a_1)} &= [0, -1] & \phi(x_2, a_2) &= [0, -2] \end{aligned} \quad (5)$$

is HLS and BBK but not WYS. To show that it is BBK, we can use the same argument used for (4). To show that it is not WYS, we notice that  $a_2$  is never optimal, so  $\text{span}\{\phi(x, a_2) \mid x \in \mathcal{X}^*(a_2)\} = \emptyset$ . This example also shows that HLS  $\cap$  BBK is not empty in general, and that BBK  $\not\subseteq$  WYS.

**WYS**  $\not\subseteq$  **BBK** The following representation:

$$\begin{array}{ll} \underline{\phi(x_1, a_1)} = [2, 0] & \phi(x_1, a_2) = [1, 0] \\ \underline{\phi(x_2, a_1)} = [0, 2] & \phi(x_2, a_2) = [0, 1] \\ \phi(x_2, a_1) = [1, 0] & \underline{\phi(x_2, a_2)} = [2, 0] \\ \phi(x_2, a_1) = [0, 1] & \underline{\phi(x_2, a_2)} = [0, 2] \end{array}$$

is WYS since each arm admits two orthogonal optimal features, but not BBK since all the features lie in the first quadrant of  $\mathbb{R}^2$ . This example also shows that WYS is not empty in general.

**BBK**  $\cap$  **HLS**  $\cap$  **WYS**  $\neq \emptyset$ . The following representation:

$$\begin{array}{ll} \underline{\phi(x_1, a_1)} = [2, 0] & \phi(x_1, a_2) = [1, 0] \\ \underline{\phi(x_2, a_1)} = [0, 2] & \phi(x_2, a_2) = [0, 1] \\ \phi(x_2, a_1) = [-2, 0] & \underline{\phi(x_2, a_2)} = [-1, 0] \\ \phi(x_2, a_1) = [0, -2] & \underline{\phi(x_2, a_2)} = [0, -1] \end{array}$$

is HLS, BBK and WYS. This can be shown using the arguments employed for the previous examples.

## B.2. Existence of HLS Representations

In this section, we prove a slightly more general version of Theorem 3.

We first prove the existence of an HLS representation in the case of finite contexts.

**Lemma 4.** *For any dimension  $d \geq 1$ , any contextual problem such that:*

- $\mathcal{X}$  is finite,
- $|\text{supp}(\rho)| \geq d$ , and
- there exists  $x \in \text{supp}(\rho)$  with  $\mu^*(x) \neq 0$ ,

*admits a  $d$ -dimensional realizable HLS representation.*

*Proof.* Consider a contextual problem  $\mathcal{P}$  with the properties stated above. Without loss of generality,  $\rho(x) > 0$  for  $x \in \{x_1, \dots, x_d\}$ , and  $\mu^*(x_1) \neq 0$ .

In this case, to prove that a  $d$ -dimensional realizable linear representation  $\phi$  is HLS, it suffices to show that  $\text{span}\{\phi^*(x_i)\}_{i=1}^d = \mathbb{R}^d$ . Indeed:

$$\lambda_{\min}(\mathbb{E}_{x \sim \rho}[\phi^*(x)\phi^*(x)^\top]) \geq \lambda_{\min}\left(\sum_{i=1}^d \rho(x_i)\phi^*(x_i)\phi^*(x_i)^\top\right) \geq \rho_{\min} \times \lambda_{\min}\left(\sum_{i=1}^d \phi^*(x_i)\phi^*(x_i)^\top\right), \quad (6)$$

where  $\rho_{\min} := \min_{i \in [d]} \rho(x_i) > 0$ . If  $\phi^*(x_1), \dots, \phi^*(x_d)$  span  $\mathbb{R}^d$ , the matrix  $\Phi$  having these vectors as rows is full rank, hence so is  $\sum_{i=1}^d \phi^*(x_i)\phi^*(x_i)^\top = \Phi^\top \Phi$ , and the minimum eigenvalue in (6) is positive.

We first consider the case  $d = 1$ . A realizable one-dimensional linear representation for  $\mathcal{P}$  is  $\phi(x, a) = \mu(x, a)$  for all  $x, a$ , with  $\theta^* = 1$ . This is HLS since  $\phi^*(x_1) = \mu^*(x_1) \neq 0$  spans  $\mathbb{R}$ .

Now consider the case  $d > 1$ . Let  $L$  be a constant such that  $L > \max_{i \in [d]} |\mu^*(x_i)|$ . Notice that  $L > |\mu^*(x_1)| > 0$ . Within the scope of this proof, we use  $\phi_i$  to denote the  $i$ -th element of vector  $\phi$ . Consider the linear representation with  $\phi_1(x, a) = \mu(x, a)$  and  $\phi_i(x, a) = L \times \mathbb{1}\{x = x_i\}$  for  $i = 2, \dots, d$  and all  $x, a$ . This is a  $d$ -dimensional linear

representation for  $\mathcal{P}$  with  $\theta^* = [1, 0, \dots, 0]^\top$ . The optimal features of the first  $d$  contexts, taken as rows, form the following matrix:

$$\begin{bmatrix} \mu^*(x_1) & 0 & 0 & \dots & 0 \\ \mu^*(x_2) & L & 0 & \dots & 0 \\ \mu^*(x_3) & 0 & L & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mu^*(x_d) & 0 & 0 & \dots & L \end{bmatrix}$$

which is strictly diagonally dominant, hence full-rank. Equivalently,  $\phi^*(x_1), \dots, \phi^*(x_d)$  span  $\mathbb{R}^d$ , completing the proof.  $\square$

Next, we prove a similar result for the continuous case.

**Lemma 5.** *For any dimension  $d \geq 1$ , any contextual problem such that:*

- $\mathcal{X}$  is a standard Borel space with measure  $\rho$  and metric  $\partial$ ,
- $x \mapsto \mu^*(x)$  is continuous on  $\mathcal{X}$ , and
- there exist  $x_1, \dots, x_d \in \text{supp}(\rho)$  with  $\mu^*(x_i) \neq 0$  for  $i = 1, \dots, d$ ,

*admits a  $d$ -dimensional realizable HLS representation.*

*Proof.* Consider a contextual problem  $\mathcal{P}$  with the properties stated above. Without loss of generality, let  $x_1, \dots, x_d \in \text{supp}(\rho)$  and  $|\mu^*(x_i)| > \epsilon_i > 0$  for  $i = 1, \dots, d$ .

We first consider the case  $d = 1$ . Define the open interval  $I_1 := (\epsilon_1, \infty)$ . Since  $\mu^*$  is continuous,  $|\mu^*|$  is also continuous. Hence, the inverse image  $E_1 = |\mu^*|^{-1}(I_1)$  is an open subset of  $\mathcal{X}$ . In particular, since  $x_1 \in E_1$ ,  $E_1$  is an open neighborhood of  $x_1$ . Since  $x_1 \in \text{supp}(\rho)$ , by definition of support<sup>10</sup>,  $\rho(E_1) > 0$ . Moreover, by definition of  $E_1$ ,  $|\mu^*(x)| > \epsilon_1$  for all  $x \in E_1$ . Consider the representation  $\phi(x, a) = \mu(x, a)$  for all  $x, a$ . This is a one-dimensional realizable linear representation for  $\mathcal{P}$  with  $\theta^* = 1$ . We have that:

$$\lambda_{\min}(\mathbb{E}_{x \sim \rho}[\phi^*(x)\phi^*(x)^\top]) = \int_{\mathcal{X}} \mu^*(x)^2 d\rho \geq \int_{E_1} \mu^*(x)^2 d\rho > \rho(E_1) \times \epsilon_1^2 > 0,$$

and the representation is HLS.

Now consider the case  $d > 1$ . For  $i = 1, \dots, d$ , define the open interval  $I_i := (\epsilon_i, \infty)$ . As in the scalar case,  $E_i = |\mu^*|^{-1}(I_i)$  is an open neighborhood of  $x_i$  of positive measure such that  $|\mu^*(x)| > \epsilon_i$  for all  $x \in E_i$ . However, these neighborhoods may not be disjoint.

To fix this, we will use the fact that  $\mathcal{X}$  is a metric space. Let  $r_0 = \min\{\partial(x_i, x_j) \mid i, j = 1, \dots, d\}$  be the minimum distance between any of the  $d$  points. Also, let  $r_i$  be the radius of the smallest open ball contained in  $E_i$ . Let  $r = \min\{r_0, r_1, \dots, r_d\}$ . Finally, let  $F_i = \mathcal{B}_{x_i, r} = \{x \in \mathcal{X} \mid \partial(x, x_i) < r\}$  be the open ball of radius  $r$  centered in  $x_i$ . The  $F_1, \dots, F_d$  defined in this way are *disjoint* open neighborhoods of  $x_1, \dots, x_d$  of positive measure such that, for all  $x \in F_i$ ,  $|\mu^*(x)| > \epsilon_i$ .

Let  $L$  be a constant such that  $L > \max_{i \in [d]} \epsilon_i$ . Notice that  $L > 0$ . Within the scope of this proof, we use  $\phi_i$  to denote the  $i$ -th element of vector  $\phi$ . Consider the linear representation with  $\phi_1(x, a) = \mu(x, a)$  and  $\phi_i(x, a) = L \times \mathbb{1}\{x \in F_i\}$  for  $i = 2, \dots, d$  and all  $x, a$ . This is a  $d$ -dimensional realizable linear representation for  $\mathcal{P}$ , with  $\theta^* = [1, 0, \dots, 0]^\top$ . Notice that, for all  $x \in F_i$ ,  $\|\phi^*(x)\|^2 = \mu^*(x)^2 + L^2 > \epsilon_i^2 + L^2$ . For  $i = 1, \dots, d$ , define vector  $\mathbf{u}(i)$  as follows:  $\mathbf{u}_1(i) = \epsilon_i$  and  $\mathbf{u}_j(i) = L \times \mathbb{1}\{x \in F_j\}$  for  $j = 1, \dots, d$ . For any vector  $\mathbf{u} \in \mathbb{R}^d$ , the only nonzero eigenvalue of  $\mathbf{u}\mathbf{u}^\top$  is  $\|\mathbf{u}\|^2$ . So if

<sup>10</sup>In a Borel space, the support of the measure is defined as the smallest set for which every open neighbourhood of every point in the set has positive measure.

$x \in F_i$ ,  $\phi^*(x)\phi^*(x)^\top \succ \mathbf{u}(i)\mathbf{u}(i)^\top$ . Finally, we have that:

$$\begin{aligned}
 \lambda_{\min} \left( \mathbb{E}_{x \sim \rho} [\phi^*(x)\phi^*(x)^\top] \right) &= \lambda_{\min} \left( \int_{\mathcal{X}} \phi^*(x)\phi^*(x)^\top d\rho \right) \\
 &\geq \lambda_{\min} \left( \sum_{i=1}^d \int_{F_i} \phi^*(x)\phi^*(x)^\top d\rho \right) \\
 &\geq \lambda_{\min} \left( \sum_{i=1}^d \int_{F_i} \mathbf{u}(i)\mathbf{u}(i)^\top d\rho \right) \\
 &= \lambda_{\min} \left( \sum_{i=1}^d \rho(F_i) \mathbf{u}(i)\mathbf{u}(i)^\top \right) \\
 &\geq \rho_{\min} \times \lambda_{\min} \left( \sum_{i=1}^d \mathbf{u}(i)\mathbf{u}(i)^\top \right),
 \end{aligned} \tag{7}$$

where  $\rho_{\min} = \min_{i \in [d]} \rho(F_i) > 0$ . In (7), we used the fact that the  $F_i$  are disjoint. It remains to show that the matrix  $U$  obtained by taking  $\mathbf{u}(1), \dots, \mathbf{u}(d)$  as rows is full-rank. But this yields the following matrix:

$$\begin{bmatrix} \epsilon_1 & 0 & 0 & \dots & 0 \\ \epsilon_2 & L & 0 & \dots & 0 \\ \epsilon_3 & 0 & L & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \epsilon_d & 0 & 0 & \dots & L \end{bmatrix} \tag{9}$$

which, by definition of  $L$ , is strictly diagonally dominant, hence full-rank. So, the matrix  $U^\top U = \sum_{i=1}^d \mathbf{u}(i)\mathbf{u}(i)^\top$  in (8) also is full-rank, its minimum eigenvalue is positive, and  $\phi$  is HLS.  $\square$

The following lemma implies that, once there exist an HLS representation for a problem, there exist infinite equivalent HLS representations (at least one for each  $d \times d$  invertible matrix):

**Lemma 6.** *Any invertible linear transformation of an HLS representation yields an HLS representation.*

*Proof.* Let  $\phi : \mathcal{X} \times [K] \rightarrow \mathbb{R}^d$  be a realizable linear representation for a problem  $\mathcal{P}$  with context space  $\mathcal{X}$  and  $K$  arms, with corresponding parameter  $\theta^*$ . Any invertible matrix  $A \in \mathbb{R}^{d \times d}$  defines an invertible linear transformation  $T_A$  in the following sense:

$$T_A(\phi)(x, a) = A^\top \phi(x, a). \tag{10}$$

First, note that  $\tilde{\phi} = T_A(\phi)$  is still a ( $d$ -dimensional) realizable linear representation for the same problem with parameter  $\tilde{\theta}^* = A^{-1}\theta^*$ , since for all  $x, a$ :

$$\mu(x, a) = \phi(x, a)^\top \theta^* = \phi(x, a)^\top A A^{-1} \theta^* = \tilde{\phi}(x, a)^\top \tilde{\theta}^*. \tag{11}$$

Now note that:

$$\lambda_{\min} \left( \mathbb{E}_{x \sim \rho} [\tilde{\phi}^*(x)\tilde{\phi}^*(x)^\top] \right) = \lambda_{\min} \left( \mathbb{E}_{x \sim \rho} [A^\top \phi^*(x)\phi^*(x)^\top A] \right).$$

If the matrix in the LHS is full-rank, so is the matrix in the RHS, because  $A$  is full-rank and the product of full-rank square matrices is always full-rank. Hence if  $\phi$  is HLS, so is  $\tilde{\phi}$ .  $\square$

These results are summarized in Lemma 3 in the paper. For simplicity, we assumed that  $\mu^*(x) \neq 0$  for all  $x \in \mathcal{X}$ . In the continuous case, we also assumed that  $\rho$  is non-degenerate, which means  $\text{supp}(\rho)$  contains an open ball. This two assumptions together guarantee there exist  $d$  (in fact, infinite) points in the support of  $\rho$  with nonzero optimal reward.

Finally, we can remove all assumptions on nonzero optimal rewards by introducing a concept of equivalence between contextual problems. We say two reward functions  $\mu$  and  $\mu'$  are equivalent provided  $\mu'(x, a) = \mu(x, a) + C$  for all  $x, a$



and some  $C \in \mathbb{R}$ . We say two contextual problems are equivalent if they have the same context space, the same arm set, the same context distribution, and equivalent reward functions. So, if the assumptions on nonzero optimal rewards from Lemma 3, 4, or 5 are not satisfied, we can always find an *equivalent* problem that admits an HLS representation by adding an appropriate constant offset to all rewards.

### B.3. Additional Properties of HLS Representations

In this section, we prove additional properties of HLS representations that will be useful to construct illustrative representation-selection problems (App. G.1).

**Lemma 7.** *Under the same assumptions of Lemma 4, for every  $d > 1$ , every contextual problem, assuming non-optimal features span  $\mathbb{R}^d$ , admits an infinite number of  $d$ -dimensional non-redundant representations that are not HLS.*

*Proof.* Lemma 4 shows that every contextual problem admits a  $d$ -dimensional HLS (hence, non-redundant) representation. We will show how to turn an HLS representation  $\phi$  of dimension  $d > 1$  into an equivalent representation that is still non-redundant but not HLS, having  $\text{rank}(M_\phi) \leq k < d$  for a  $k$  of choice, where  $M_\phi = \mathbb{E}_{x \sim \rho}[\phi^*(x)\phi^*(x)^\top]$ . As shown in the proof of Lemma 6, an invertible linear transformation does not affect the rank of  $M_\phi$ . So, one can easily obtain an infinite number of equivalent representations with the same property.

Let  $\Phi^* \in \mathbb{R}^{N \times d}$  be a matrix having as rows the optimal features, after removing contexts that are not in the support of  $\rho$ . Let  $q = N - k + 1$ , where  $k$  is the desired rank, and  $\Phi_q^*$  denotes  $\Phi^*$  with all the  $d$  columns but only the first  $q$  rows. We will only modify  $\Phi_q^*$ , leaving all other features unchanged. Let  $\mu^* \in \mathbb{R}^N$  be the vector of optimal rewards, and  $\mu_q^*$  be the sub-array of its first  $q$  elements. The modified features are:

$$\tilde{\Phi}_q^* = \frac{\mu_q^*(\mu_q^*)^\top}{\|\mu_q^*\|^2} \Phi_q^*.$$

The new representation  $\tilde{\phi}$  obtained in this way is equivalent to  $\phi$  (with  $\tilde{\theta}^* = \theta^*$ ) since:

$$\tilde{\Phi}_q^* \theta^* = \frac{\mu_q^*(\mu_q^*)^\top}{\|\mu_q^*\|^2} \Phi_q^* \theta^* = \frac{\mu_q^*(\mu_q^*)^\top \mu_q^*}{\|\mu_q^*\|^2} = \mu_q^* = \Phi_q^* \theta^*,$$

and all other features are unchanged. However, note that  $\tilde{\Phi}_q^*$  has rank one, since it is obtained by multiplication with the rank-one matrix  $\mu_q^*(\mu_q^*)^\top$ . This means that the first  $q = N - k + 1$  optimal features of  $\tilde{\phi}$  are all linearly dependent. With the remaining  $k - 1$  features, the rank of  $\tilde{\Phi}^*$  is at most  $k < d$ , so  $\tilde{\phi}$  is not HLS. Since we only modified optimal features, assuming non-optimal features of  $\phi$  span  $\mathbb{R}^d$ ,  $\tilde{\phi}$  is still non-redundant.  $\square$

**Lemma 8.** *Let  $\Phi \in \mathbb{R}^{N \times K \times d}$  be a random matrix whose elements are sampled i.i.d. from a non-degenerate distribution (e.g., a standard normal). Let  $\theta^* \in \mathbb{R}^d$  be any vector and consider the contextual problem with context set  $[N]$ , action set  $[K]$ , uniform context distribution  $\rho$ , and  $\mu : (x, a) \mapsto \Phi[xK + a]^\top \theta^*$ , where  $[\cdot]$  selects rows. Then, almost surely, the representation  $\phi : (x, a) \mapsto \Phi[xK + a]$  is HLS for this problem.*

*Proof.* Let  $\Phi^*$  be the submatrix of  $\Phi$  with only rows  $xK + a$  such that  $a$  is an optimal action for  $x$ . Notice that  $\phi$  is HLS if and only if  $\Phi^*$  is full-rank. But the rows of  $\Phi^*$  are sampled i.i.d. from a non degenerate  $d$ -variate distribution (e.g., a  $d$ -variate Gaussian with positive variance in all directions). So,  $\Phi^*$  is full-rank almost surely (Eaton & Perlman, 1973).  $\square$

## C. Constant Regret with a HLS Representation

### C.1. Preliminary Results

**Proposition 3** (Abbasi-Yadkori et al., 2011). *Consider any linear contextual bandit problem with noise standard deviation  $\sigma$ . Run LINUCB with confidence parameter  $\delta$ , regularization parameter  $\lambda$ , and a  $d_\phi$ -dimensional realizable linear representation such that  $\sup_{x,a} \|\phi(x, a)\|_2 \leq L_\phi$ , and  $\|\theta_\phi^*\|_2 \leq S_\phi$ . The following good event holds with probability at least  $1 - \delta$ :*

$$\mathcal{G}_\phi(\delta) := \left\{ \forall t \geq 1, \|\theta_{t\phi} - \theta_\phi^*\|_{V_{t\phi}} \leq \beta_{t\phi}(\delta) \right\}, \quad (12)$$

where

$$\beta_{t\phi}(\delta) := \sigma \sqrt{2 \ln \left( \frac{\det(V_{t\phi})^{1/2} \det(\lambda I_{d_\phi})^{-1/2}}{\delta} \right)} + \sqrt{\lambda} S_\phi \quad (13)$$

$$\leq \sigma \sqrt{2 \log(1/\delta) + d_\phi \log(1 + (t-1)L_\phi^2/(\lambda d_\phi))} + \sqrt{\lambda} S_\phi. \quad (14)$$

**Proposition 4** (Abbasi-Yadkori et al., 2011). *Under the same assumptions of Proposition 3, assuming the good event  $\mathcal{G}_\phi(\delta)$  holds, for all  $t \geq 1$ , the instantaneous regret of LINUCB is bounded as:*

$$r_t \leq 2\beta_{t\phi}(\delta) \|\phi(x_t, a_t)\|_{V_{t\phi}^{-1}}.$$

### Proof of Proposition 1.

*Proof.* This is a variant of Theorem 5 by Abbasi-Yadkori et al. (2011) where we do not assume the existence of a unique optimal feature vector over  $\mathcal{X}$ . The proof is easier, and the result has a worse dependence on the feature dimension.

The assumptions of this proposition are the same as Prop. 3. We will show that the logarithmic regret bound holds under the good event  $\mathcal{G}_\phi(\delta)$ , which in turn holds with probability at least  $1 - \delta$ . Here and in the rest of the paper, we further assume that  $\Delta_{\max} \geq 1$ ,  $S_\phi \geq 1$ ,  $\lambda \geq 1$ , and  $\sigma \geq 1$ . These are technical assumptions that can be easily removed by properly clipping the constants in the regret bound.

From Prop. 4 and  $\beta_{t\phi}(\delta) \geq 1$ :

$$\begin{aligned} r_t &\leq 2\beta_{t\phi}(\delta) \|\phi(x_t, a_t)\|_{V_{t\phi}^{-1}} \\ &\leq \min \left\{ 2\beta_{t\phi}(\delta) \|\phi(x_t, a_t)\|_{V_{t\phi}^{-1}}, \Delta_{\max} \right\} \\ &\leq 2\Delta_{\max} \beta_{t\phi}(\delta) \min \left\{ \|\phi(x_t, a_t)\|_{V_{t\phi}^{-1}}, 1 \right\}. \end{aligned}$$

From the fact that  $\beta_{t\phi}(\delta)$  is increasing and the Elliptical Potential Lemma (e.g., Abbasi-Yadkori et al., 2011, Lemma 11):

$$\begin{aligned} \sum_{k=1}^t r_k^2 &\leq 4\Delta_{\max}^2 \beta_{t\phi}(\delta)^2 \sum_{k=1}^t \min \left\{ \|\phi(x_k, a_k)\|_{V_{k\phi}^{-1}}^2, 1 \right\} \\ &\leq 8\Delta_{\max}^2 \beta_{t\phi}(\delta)^2 d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)). \end{aligned}$$

Since either  $r_t \geq \Delta$  or  $r_t = 0$ :

$$\begin{aligned} R_t &= \sum_{k=1}^t r_k \leq \sum_{k=1}^t \frac{r_k^2}{\Delta} \leq \frac{8\Delta_{\max}^2 \beta_{t\phi}(\delta)^2 d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi))}{\Delta} \\ &\leq \frac{8\Delta_{\max}^2 \left( \sigma \sqrt{2 \log(1/\delta) + d_\phi \log(1 + (t-1)L_\phi^2/(\lambda d_\phi))} + \sqrt{\lambda} S_\phi \right)^2 d_\phi \log(1 + L_\phi^2/(\lambda d_\phi))}{\Delta} \\ &\leq \frac{32\Delta_{\max}^2 \lambda S_\phi^2 \sigma^2 \left( 2 \log(1/\delta) + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)) \right)^2}{\Delta}. \end{aligned}$$

□

**Proposition 5** (Abbasi-Yadkori et al., 2011). *Under the same assumptions of Proposition 3, assuming the good event  $\mathcal{G}_\phi(\delta)$  holds, for all  $t \geq 1$ , the number of suboptimal pulls of LINUCB up to time  $t$  is at most:*

$$g_{t\phi}(\delta) := \frac{32\Delta_{\max}^2 \lambda S_\phi^2 \sigma^2 \left( 2 \log(1/\delta) + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)) \right)^2}{\Delta^2}.$$

*Proof.* A similar bound can be found in the proof of Theorem 5 by Abbasi-Yadkori et al. (2011). Again, we do not assume a unique optimal feature vector.

Let  $n_t$  be the number of suboptimal pulls up to time  $t$ . Since  $R_t \geq n_t \Delta$ :

$$n_t \leq \frac{R_t}{\Delta}.$$

The regret bound from Prop. 1, which holds under  $\mathcal{G}_\phi(\delta)$ , completes the proof.  $\square$

**Lemma 9.** *Under the same assumptions of Proposition 3, assuming the good event  $\mathcal{G}_\phi(\delta)$  holds, with probability  $1 - \delta$ , for all  $t \geq 1$ :*

$$\lambda_{\min}(V_{t+1,\phi}) \geq \lambda + t\lambda_{\phi,\text{HLS}} - 8L_\phi^2 \sqrt{t \log(2d_\phi t / \delta)} - L_\phi^2 g_{t\phi}(\delta), \quad (15)$$

where  $g_{t\phi}(\delta)$  is from Prop. 5.

*Proof.* Since  $V_{t+1,\phi}$  is symmetric:

$$\begin{aligned} V_{t+1,\phi} &= \lambda I_{d_\phi} + \sum_{k=1}^t \phi(x_k, a_k) \phi(x_k, a_k)^\top \\ &= \lambda I_{d_\phi} + \sum_{k=1}^t \mathbb{1}\{a_k = a_{x_k}^*\} \phi(x_k, a_k) \phi(x_k, a_k)^\top + \sum_{k=1}^t \mathbb{1}\{a_k \neq a_{x_k}^*\} \phi(x_k, a_k) \phi(x_k, a_k)^\top \\ &\succeq \lambda I_{d_\phi} + \sum_{k=1}^t \mathbb{1}\{a_k = a_{x_k}^*\} \phi^*(x_k) \phi^*(x_k)^\top \\ &= \lambda I_{d_\phi} + \sum_{k=1}^t \phi^*(x_k) \phi^*(x_k)^\top - \sum_{k=1}^t \mathbb{1}\{a_k \neq a_{x_k}^*\} \phi^*(x_k) \phi^*(x_k)^\top \\ &\succeq \lambda I_{d_\phi} + \sum_{k=1}^t \phi^*(x_k) \phi^*(x_k)^\top - g_{t\phi}(\delta) L_\phi^2 I_{d_\phi} \\ &= \lambda I_{d_\phi} + t \mathbb{E}_{x \sim \rho} [\phi^*(x) \phi^*(x)^\top] - \sum_{k=1}^t (\mathbb{E}_{x \sim \rho} [\phi^*(x) \phi^*(x)^\top] - \phi^*(x_k) \phi^*(x_k)^\top) - L_\phi^2 g_{t\phi}(\delta) I_{d_\phi} \\ &= \lambda I_{d_\phi} + t \mathbb{E}_{x \sim \rho} [\phi^*(x) \phi^*(x)^\top] - \sum_{k=1}^t X_k - L_\phi^2 g_{t\phi}(\delta) I_{d_\phi}, \end{aligned} \quad (16)$$

where  $g_{t\phi}(\delta)$  is the upper bound on suboptimal pulls from Prop. 5, which holds under the good event  $\mathcal{G}_\phi(\delta)$ , and  $X_k = \mathbb{E}_{x \sim \rho} [\phi^*(x) \phi^*(x)^\top] - \phi^*(x_k) \phi^*(x_k)^\top$ . Since the matrix in (26) is still symmetric, by definition of  $\lambda_{\phi,\text{HLS}}$ :

$$\lambda_{\min}(V_{t+1,\phi}) \geq \lambda + t\lambda_{\phi,\text{HLS}} - \lambda_{\max} \left( \sum_{k=1}^t X_k \right) - L_\phi^2 g_{t\phi}(\delta).$$

We bound the third term using a matrix Azuma inequality by Tropp (2012). First, notice that  $\mathbb{E}_k[X_k] = 0$ . Also, since  $X_k$  is symmetric:

$$X_k^2 \preceq \lambda_{\max}(X_k^2) I_{d_\phi} \preceq \|X_k\|^2 I_{d_\phi} \preceq 4L_\phi^4 I_{d_\phi}.$$

Hence, from Prop. 27, with probability at least  $1 - \delta'_t$ , for all  $t \geq 1$ :<sup>11</sup>

$$\lambda_{\max} \left( \sum_{k=1}^t X_k \right) \leq 4L_\phi^2 \sqrt{2t \log(d_\phi / \delta'_t)}.$$

<sup>11</sup>Notice that this is true regardless of  $\mathcal{G}_\phi(\delta)$ . This will be useful to bound expected regret in App. C.3.

We set  $\delta'_t = \delta/(2t^2)$  and perform a union bound over time. Finally, with probability at least  $1 - \delta$ , for all  $t \geq 1$ :

$$\lambda_{\max} \left( \sum_{k=1}^t X_k \right) \leq 4L_\phi^2 \sqrt{2t \log(2d_\phi t^2/\delta)} \leq 8L_\phi^2 \sqrt{t \log(2d_\phi t/\delta)}.$$

□

## C.2. High-Probability Regret Bound

**Lemma 10.** *Consider a contextual bandit problem with realizable linear representation  $\phi$  satisfying the HLS condition. Assume  $\Delta > 0$ ,  $\max_{x,a} \|\phi(x,a)\|_2 \leq L_\phi$  and  $\|\theta_\phi^*\|_2 \leq S_\phi$ . Then, assuming the good event  $\mathcal{G}_\phi(\delta)$  holds, with probability at least  $1 - \delta$ , there exists a (constant) time  $\tau_\phi$  such that, for all  $t \geq \tau_\phi$ , the instantaneous regret of LINUCB run with confidence parameter  $\delta$  is  $r_{t+1} = 0$ .*

*Proof.* From Prop. 4, under the good event  $\mathcal{G}_\phi(\delta)$ :

$$\begin{aligned} r_{t+1} &\leq 2\beta_{t+1,\phi}(\delta) \|\phi(x_{t+1}, a_{t+1})\|_{V_{t+1,\phi}^{-1}} \\ &\leq 2\beta_{t+1,\phi}(\delta) L_\phi \sqrt{\lambda_{\max}(V_{t+1,\phi}^{-1})} \\ &\leq 2\beta_{t+1,\phi}(\delta) \frac{L_\phi}{\sqrt{\lambda_{\min}(V_{t+1,\phi})}}. \end{aligned}$$

Since either  $r_{t+1} \geq \Delta$  or  $r_{t+1} = 0$ , we just need to show that, for all  $t \geq \tau_\phi$ :

$$2\beta_{t+1,\phi}(\delta) \frac{L_\phi}{\sqrt{\lambda_{\min}(V_{t+1,\phi})}} < \Delta.$$

From Lemma 9, which holds with probability at least  $1 - \delta$  under  $\mathcal{G}_\phi(\delta)$ , rearranging:

$$t\lambda_{\phi,\text{HLS}} > \frac{4\beta_{t+1,\phi}(\delta)^2 L_\phi^2}{\Delta^2} + 8L_\phi^2 \sqrt{t \log(2d_\phi t/\delta)} + L_\phi^2 g_{t\phi}(\delta) - \lambda.$$

By replacing  $\beta_{t+1,\phi}(\delta)$  with its bound from 14 and  $g_{t\phi}(\delta)$  with its definition from Prop. 5 a sufficient condition is:

$$\begin{aligned} t\lambda_{\phi,\text{HLS}} &> \frac{4 \left( \sigma \sqrt{2 \log(1/\delta)} + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)) + \sqrt{\lambda} S_\phi \right)^2 L_\phi^2}{\Delta^2} + 8L_\phi^2 \sqrt{t \log(2d_\phi t/\delta)} \\ &\quad + \frac{32\Delta_{\max}^2 \lambda L_\phi^2 S_\phi^2 \sigma^2 \left( 2 \log(1/\delta) + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)) \right)^2}{\Delta^2} - \lambda. \end{aligned} \quad (17)$$

Since  $\phi$  is HLS,  $\lambda_{\phi,\text{HLS}} > 0$  and the LHS is linear in  $t$ , while the RHS is sublinear. This means we can find a sufficiently large constant  $\tau_\phi$  such that (17) holds for all  $t \geq \tau_\phi$ . One such time is derived explicitly in Appendix C.6. □

### Proof of Lemma 2.

*Proof.* First assume the good event  $\mathcal{G}_\phi(\delta)$  holds. From Lemma 10, with probability  $1 - \delta$ , the instantaneous regret is zero after  $\tau_\phi$ . So, we can replace  $n$  with  $\tau_\phi$  in the anytime regret upper bound from Prop. 1, which also holds under  $\mathcal{G}_\phi(\delta)$ . More precisely, if  $n \geq \tau_\phi$ :

$$R_n = \sum_{t=1}^n r_t = \sum_{t=1}^{\tau_\phi} r_t + \underbrace{\sum_{t=\tau_\phi+1}^n r_t}_0 = \sum_{t=1}^{\tau_\phi} r_t.$$

So, in any case:

$$R_n \leq \sum_{t=1}^{\min\{n, \tau_\phi\}} r_t = R_{\min\{n, \tau_\phi\}},$$



which can be bounded with Prop. 1 to obtain:

$$R_n \leq \frac{32\Delta_{\max}^2 \lambda S_\phi^2 \sigma^2}{\Delta} \left( 2 \ln \left( \frac{1}{\delta} \right) + d_\phi \ln \left( 1 + \frac{\min\{\tau_\phi, n\} L_\phi^2}{\lambda d_\phi} \right) \right)^2.$$

From Prop. 3, the good event may fail with probability at most  $\delta$ . Lemma 10 may fail anyway with probability at most  $\delta$ . From a union bound, the overall failure probability is at most  $2\delta$ .  $\square$

### C.3. Expected Regret Bound

To establish an expected-regret guarantee, we need to consider a slight variant of LINUCB that employs an adaptive confidence parameter schedule  $(\delta_t)_{t=1}^\infty$ . For now, we just assume  $\delta_t \in (0, 1)$  for all  $t$  and is a decreasing function of  $t$ .

We will need the following good events:

- $\mathcal{G}_{t\phi} := \left\{ \left\| \theta_{t\phi} - \theta_\phi^* \right\|_{V_{t\phi}} \leq \beta_{t\phi}(\delta_t) \right\}$ . From (Abbasi-Yadkori et al., 2011) we know that this holds with probability at least  $1 - \delta_t$ . Under this event, we have the instantaneous regret upper bound from Prop. 4 restricted to time  $t$ .
- $\mathcal{F}_{t\phi} := \{n_t \leq g_{t\phi}(\delta_t)\}$ , where  $n_t$  is the number of suboptimal pulls of LINUCB up to time  $t$  and  $g_{t\phi}$  is defined in Prop. 5. This event holds if  $\mathcal{G}_{k\phi}$  holds for all  $k \leq t$ . Indeed:

$$\begin{aligned} n_t &\leq \frac{R_t}{\Delta} \leq \sum_{k=1}^t \frac{r_k^2}{\Delta} \\ &\leq \sum_{k=1}^t \frac{4\Delta_{\max}^2 \beta_{k\phi}(\delta_k)^2 \min \left\{ \|\phi(x_k, a_k)\|_{V_{k\phi}^{-1}}^2, 1 \right\}}{\Delta} \\ &\leq 4\Delta_{\max}^2 \beta_{t\phi}(\delta_t)^2 \sum_{k=1}^t \frac{\min \left\{ \|\phi(x_k, a_k)\|_{V_{k\phi}^{-1}}^2, 1 \right\}}{\Delta} \\ &\leq g_{t\phi}(\delta_t), \end{aligned} \tag{18}$$

where (18) uses the fact that  $\delta_t$  is decreasing in  $t$  and  $\beta_{t\phi}(\delta)$  is increasing in  $t$  and decreasing in  $\delta$ , and the last inequality uses the same algebraic manipulations of the proof of Prop. 5. Hence, from a union bound,  $\mathbb{P}(\mathcal{F}_{t\phi}) \geq 1 - \sum_{k=1}^t \delta_k$ .

- $\mathcal{E}_{t\phi} := \left\{ \lambda_{\max} \left( \sum_{k=1}^t \mathbb{E}_{x \sim \rho} [\phi^*(x) \phi^*(x)^\top] - \phi^*(x_k) \phi^*(x_k)^\top \right) \leq 4L_\phi^2 \sqrt{2t \log(d_\phi / \delta'_t)} \right\}$ . This is the event we used to concentrate contexts in the proof of Lemma 9. From Prop. 27, it holds with probability  $\delta'_t$ . Since it is independent from the behavior of LINUCB, we can set  $\delta'_t$  to any convenient value regardless of the confidence schedule  $\delta_t$ .

**Lemma 11.** *Consider a contextual bandit problem with realizable linear representation  $\phi$  satisfying the HLS condition. Assume  $\Delta > 0$ ,  $\max_{x,a} \|\phi(x, a)\|_2 \leq L$  and  $\|\theta_\phi^*\|_2 \leq S$ . Then, the expected regret of OFUL run with adaptive confidence schedule  $\delta_t = 1/t^3$  after  $n \geq 1$  steps is at most:*

$$\mathbb{E} R_n \leq \frac{32\Delta_{\max}^2 \lambda S_\phi^2 \sigma^2}{\Delta} \left( 6 \ln(\min\{\tilde{\tau}_\phi, n\}) + d_\phi \ln \left( 1 + \frac{\min\{\tilde{\tau}_\phi, n\} L_\phi^2}{\lambda d_\phi} \right) \right)^2 + 26,$$

where  $\tilde{\tau}_\phi$  is a constant independent from  $n$ .

*Proof.* Fix a time  $t$  and assume the three good events  $\mathcal{G}_{t+1,\phi}$ ,  $\mathcal{F}_{t\phi}$ , and  $\mathcal{E}_{t\phi}$  hold. Following the proof of Lemma 10 we can show that, thanks to  $\mathcal{G}_{t+1,\phi}$ :

$$r_{t+1} \leq 2\beta_{t+1,\phi}(\delta_{t+1}) \frac{L_\phi}{\sqrt{\lambda_{\min}(V_{t+1,\phi})}}.$$

Following the proof of Lemma 9 we can show that, thanks to  $\mathcal{G}_{t+1,\phi}$  and  $\mathcal{E}_{t\phi}$ :

$$\lambda_{\min}(V_{t+1,\phi}) \geq \lambda + t\lambda_{\phi,\text{HLS}} - 4L_\phi^2 \sqrt{2t \log(d_\phi/\delta'_t)} - L_\phi^2 n_t,$$

and  $n_t \leq g_{t\phi}(\delta_t)$  thanks to  $\mathcal{F}_{t\phi}$ . Putting all together, as observed in the proof of Lemma 10, a sufficient condition for  $r_{t+1} = 0$  is:

$$t\lambda_{\phi,\text{HLS}} > \frac{4\beta_{t+1,\phi}(\delta_{t+1})^2 L_\phi^2}{\Delta^2} + 4L_\phi^2 \sqrt{2t \log(d_\phi/\delta'_t)} + L_\phi^2 g_{t\phi}(\delta_t) - \lambda.$$

With our choice of  $\delta_t$ , by setting  $\delta'_t = 1/t^2$ , this becomes:

$$\begin{aligned} t\lambda_{\phi,\text{HLS}} > \frac{4 \left( \sigma \sqrt{6 \log(t+1) + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi))} + \sqrt{\lambda} S_\phi \right)^2 L_\phi^2}{\Delta^2} + 8L_\phi^2 \sqrt{2t \log(d_\phi t)} \\ + \frac{32\Delta_{\max}^2 \lambda L_\phi^2 S_\phi^2 \sigma^2 \left( 6 \log(t) + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)) \right)^2}{\Delta^2} - \lambda. \end{aligned} \quad (19)$$

The RHS is still sublinear in  $t$ , so, if  $\phi$  is HLS, we can find a sufficiently large constant  $\tilde{\tau}_\phi$  such that (19) is always satisfied for  $t \geq \tilde{\tau}_\phi$ .

If  $n \geq \tilde{\tau}_\phi$ , we can decompose the expected regret as follows:

$$\mathbb{E} R_n \leq \underbrace{\mathbb{E} \sum_{t=1}^{\tilde{\tau}_\phi} r_t}_{(a)} + \underbrace{\mathbb{E} \sum_{t=\tilde{\tau}_\phi+1}^n r_t}_{(b)}. \quad (20)$$

The first summation is just the expected regret of LINUCB up to time  $\tilde{\tau}_\phi$ :

$$\begin{aligned} (a) &= \mathbb{E} \sum_{t=1}^{\tilde{\tau}_\phi} \mathbb{1}\{\mathcal{G}_{t\phi}\} r_t + \mathbb{E} \sum_{t=1}^{\tilde{\tau}_\phi} \mathbb{1}\{\neg\mathcal{G}_{t\phi}\} r_t \\ &\leq \sum_{t=1}^{\tilde{\tau}_\phi} \beta_{k\phi}(\delta_k) \|\phi(x_k, a_k)\|_{V_k^{-1}} + 2 \sum_{t=1}^{\tilde{\tau}_\phi} \delta_t \\ &\leq \beta_{\tilde{\tau}_\phi,\phi}(\delta_{\tilde{\tau}_\phi}) \sum_{t=1}^{\tilde{\tau}_\phi} \|\phi(x_k, a_k)\|_{V_k^{-1}} + 2 \sum_{t=1}^{\infty} \delta_t \\ &\leq \frac{32\Delta_{\max}^2 \lambda S_\phi^2 \sigma^2}{\Delta} \left( 6 \ln(\tilde{\tau}_\phi) + d_\phi \ln \left( 1 + \frac{\tilde{\tau}_\phi L_\phi^2}{\lambda d_\phi} \right) \right)^2 + 2 \sum_{t=1}^{\infty} t^{-3} \\ &\leq \frac{32\Delta_{\max}^2 \lambda S_\phi^2 \sigma^2}{\Delta} \left( 6 \ln(\tilde{\tau}_\phi) + d_\phi \ln \left( 1 + \frac{\tilde{\tau}_\phi L_\phi^2}{\lambda d_\phi} \right) \right)^2 + 3, \end{aligned}$$

where we used the same algebraic manipulations used to prove Prop. 1. For the second summation:

$$\begin{aligned} (b) &= \mathbb{E} \sum_{t=\tilde{\tau}_\phi}^{n-1} r_{t+1} \\ &= \mathbb{E} \sum_{t=\tilde{\tau}_\phi}^{n-1} \underbrace{\mathbb{1}\{\mathcal{G}_{t+1,\phi} \cap \mathcal{F}_{t\phi} \cap \mathcal{E}_{t\phi}\}}_{=0} r_{t+1} + \mathbb{E} \sum_{t=\tilde{\tau}_\phi}^{n-1} \mathbb{1}\{\neg\mathcal{G}_{t+1,\phi} \cup \neg\mathcal{F}_{t\phi} \cup \neg\mathcal{E}_{t\phi}\} r_{t+1} \\ &\leq 2 \sum_{t=1}^{\infty} \left( \delta_t + \sum_{k=1}^t k^{-3} + t^{-2} \right) \\ &\leq 2 \sum_{t=1}^{\infty} t^{-3} + 2 \sum_{t=1}^{\infty} \sum_{k=1}^t \frac{1}{k^3} + 4 \leq 2 \sum_{t=1}^{\infty} \sum_{k=t}^{\infty} k^{-3} + 7 \leq 2 \sum_{t=1}^{\infty} t^{-3/2} \sum_{k=t}^{\infty} k^{-3/2} + 7 \leq 6 \sum_{t=1}^{\infty} t^{-3/2} + 7 \leq 23. \end{aligned}$$

In any case, from the same argument used to bound (a):

$$\mathbb{E} R_n \leq \frac{32\Delta_{\max}^2 \lambda S_\phi^2 \sigma^2}{\Delta} \left( 6 \ln(n) + d_\phi \ln \left( 1 + \frac{nL_\phi^2}{\lambda d_\phi} \right) \right)^2 + 3.$$

Putting everything together, we obtain our statement.  $\square$

#### C.4. Margin Condition (Proof of Theorem 1)

Before proving Th. 1, we need to generalize some of the previous results for the case where the minimum gap can be arbitrarily small but the margin condition (Asm. 1) holds. First, we show two immediate results that bound the regret suffered by LINUCB on rounds where the minimum gap is above a given value.

**Lemma 12.** *Under the good event  $G_\phi(\delta)$ , for all  $\epsilon > 0$  and  $t > 0$ , the  $t$ -step regret of LINUCB on rounds where the minimum gap is at least  $\epsilon$  can be bounded as*

$$R_t^\epsilon := \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) \geq \epsilon \} r_k \leq \frac{32\lambda(\Delta_{\max} S_\phi \sigma)^2}{\epsilon} (2 \log(1/\delta) + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)))^2. \quad (21)$$

Moreover, the number of sub-optimal pulls performed by LINUCB on rounds where the minimum gap is at least  $\epsilon$  can be bounded as

$$g_{t\phi}^\epsilon(\delta) := \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) > \epsilon, r_k > 0 \} \leq \frac{32\lambda(\Delta_{\max} S_\phi \sigma)^2}{\epsilon^2} (2 \log(1/\delta) + d_\phi \log(1 + tL_\phi^2/(\lambda d_\phi)))^2 \quad (22)$$

*Proof.* It is easy to see that

$$R_t^\epsilon = \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) \geq \epsilon \} r_k \leq \sum_{k=1}^t \mathbb{1} \{ r_k > 0, \Delta(x_k) \geq \epsilon \} \frac{r_k^2}{\epsilon} \leq \frac{1}{\epsilon} \sum_{k=1}^t r_k^2. \quad (23)$$

From here the first result follows by reproducing the proof of the regret bound in Prop. 1. The second result is immediate from

$$R_t^\epsilon = \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) \geq \epsilon, r_k > 0 \} r_k \geq \epsilon g_{t\phi}^\epsilon(\delta) \implies g_{t\phi}^\epsilon(\delta) \leq \frac{R_t^\epsilon}{\epsilon}. \quad (24)$$

$\square$

Next, we generalize Lem. 9.

**Lemma 13.** *Under Asm. 1, for any  $\epsilon > 0$  and  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ , for every  $t > 0$ ,*

$$\lambda_{\min}(V_{t+1,\phi}) \geq \lambda + t(\lambda_{\phi,\text{HLS}} - L_\phi^2 C \epsilon^\alpha) - 4(1 + C \epsilon^\alpha) L_\phi^2 \sqrt{t \log(2d_\phi t/\delta)} - L_\phi^2 g_{t\phi}^\epsilon(\delta). \quad (25)$$

*Proof.* Using a similar decomposition as in the proof of Lem. 9,

$$\begin{aligned}
 V_{t+1,\phi} &= \lambda I_{d_\phi} + \sum_{k=1}^t \phi(x_k, a_k) \phi(x_k, a_k)^\top \\
 &\succeq \lambda I_{d_\phi} + \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) \geq \epsilon, r_k = 0 \} \phi^*(x_k) \phi^*(x_k)^\top \\
 &= \lambda I_{d_\phi} + \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) \geq \epsilon \} \phi^*(x_k) \phi^*(x_k)^\top - \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) \geq \epsilon, r_k > 0 \} \phi^*(x_k) \phi^*(x_k)^\top \\
 &\succeq \lambda I_{d_\phi} + \sum_{k=1}^t \mathbb{1} \{ \Delta(x_k) \geq \epsilon \} \phi^*(x_k) \phi^*(x_k)^\top - g_{t\phi}^\epsilon(\delta) L_\phi^2 I_{d_\phi} \\
 &= \lambda I_{d_\phi} + t \mathbb{E}_{x \sim \rho} [\mathbb{1} \{ \Delta(x) \geq \epsilon \} \phi^*(x) \phi^*(x)^\top] - \sum_{k=1}^t X_k - L_\phi^2 g_{t\phi}^\epsilon(\delta) I_{d_\phi}, \tag{26}
 \end{aligned}$$

where  $g_{t\phi}^\epsilon(\delta)$  is given in Lem. 12 and  $X_k := \mathbb{E}_{x \sim \rho} [\mathbb{1} \{ \Delta(x) \geq \epsilon \} \phi^*(x) \phi^*(x)^\top] - \mathbb{1} \{ \Delta(x_k) \geq \epsilon \} \phi^*(x_k) \phi^*(x_k)^\top$ . Using the margin condition,

$$\begin{aligned}
 \lambda_{\min} (\mathbb{E}_{x \sim \rho} [\mathbb{1} \{ \Delta(x) \geq \epsilon \} \phi^*(x) \phi^*(x)^\top]) &\geq \lambda_{\phi, \text{HLS}} - \lambda_{\max} (\mathbb{E}_{x \sim \rho} [\mathbb{1} \{ \Delta(x) < \epsilon \} \phi^*(x) \phi^*(x)^\top]) \\
 &\geq \lambda_{\phi, \text{HLS}} - L_\phi^2 \mathbb{P}(\Delta(x) < \epsilon) \geq \lambda_{\phi, \text{HLS}} - L_\phi^2 C \epsilon^\alpha.
 \end{aligned}$$

Thus,

$$\lambda_{\min}(V_{t+1,\phi}) \geq \lambda + t(\lambda_{\phi, \text{HLS}} - L_\phi^2 C \epsilon^\alpha) - \lambda_{\max} \left( \sum_{k=1}^t X_k \right) - L_\phi^2 g_{t\phi}^\epsilon(\delta).$$

As in Lem. 9, we can bound the third term using a matrix Azuma inequality by Tropp (2012). Note that  $\mathbb{E}_k[X_k] = 0$  and  $\|X_k\| \leq (1 + \mathbb{P}(\Delta(x) < \epsilon)) L_\phi^2 \leq (1 + C \epsilon^\alpha) L_\phi^2$ . Hence, from Prop. 27, with probability at least  $1 - \delta'_t$ , for all  $t \geq 1$ :

$$\lambda_{\max} \left( \sum_{k=1}^t X_k \right) \leq 2(1 + C \epsilon^\alpha) L_\phi^2 \sqrt{t \log(d_\phi / \delta'_t)}.$$

Setting  $\delta'_t = \delta / (2t^2)$  and taking a union bound over time, we have that, with probability at least  $1 - \delta$ , for all  $t \geq 1$ ,

$$\lambda_{\max} \left( \sum_{k=1}^t X_k \right) \leq 2(1 + C \epsilon^\alpha) L_\phi^2 \sqrt{t \log(2d_\phi t^2 / \delta)} \leq 4(1 + C \epsilon^\alpha) L_\phi^2 \sqrt{t \log(2d_\phi t / \delta)}.$$

□

Finally, we need the following upper bound on the regret suffered by LINUCB on (possibly random) subsets of rounds.

**Lemma 14** (Cf. the proof of Lemma 13 in Tirinzoni et al. (2020)). *Let  $\{E_t\}_{t \geq 1}$  be any sequence of events. Then, under the good event  $G_\phi(\delta)$ , the  $n$ -step regret of LINUCB on rounds where the corresponding event holds can be upper bounded by*

$$\sum_{t=1}^n \mathbb{1} \{ E_t \} r_t \leq \sqrt{8 \Delta_{\max}^2 N_n d_\phi \log(1 + N_n L_\phi^2 / (\lambda d_\phi))} \left( \sqrt{\lambda} S_\phi + \sigma \sqrt{2 \log(1/\delta) + d_\phi \log(1 + n L_\phi^2 / (\lambda d_\phi))} \right), \tag{27}$$

where  $N_n := \sum_{t=1}^n \mathbb{1} \{ E_t \}$ .

*Proof.* Using the standard regret decomposition for LINUCB,

$$\begin{aligned} \sum_{t=1}^n \mathbb{1}\{E_t\} r_t &\leq \sqrt{N_n \sum_{t=1}^n \mathbb{1}\{E_t\} r_t^2} \leq \sqrt{4N_n \Delta_{\max}^2 \beta_{n\phi}^2(\delta) \sum_{t=1}^n \mathbb{1}\{E_t\} \min\{\|\phi(x_t, a_t)\|_{V_{t\phi}^{-1}}^2, 1\}} \\ &\leq \sqrt{4N_n \Delta_{\max}^2 \beta_{n\phi}^2(\delta) \sum_{t=1}^n \mathbb{1}\{E_t\} \min\{\|\phi(x_t, a_t)\|_{\tilde{V}_{t\phi}^{-1}}^2, 1\}} \\ &\leq \sqrt{8N_n \Delta_{\max}^2 \beta_{n\phi}^2(\delta) d_\phi \log(1 + N_n L_\phi^2 / (\lambda d_\phi))}, \end{aligned}$$

where  $\tilde{V}_{t\phi} := \lambda + \sum_{k=1}^t \mathbb{1}\{E_k\} \phi(x_k, a_k) \phi(x_k, a_k)^T \preceq V_{t\phi}$  since  $V_{t\phi} - \tilde{V}_{t\phi} \succeq 0$ . This implies  $\tilde{V}_{t\phi}^{-1} \succeq V_{t\phi}^{-1}$ , which implies the norm inequality. The last inequality is from the elliptical potential lemma (Abbasi-Yadkori et al., 2011).  $\square$

**Proof of Theorem 1.** We begin by proving the regret bound for LINUCB with margin condition and without a HLS representation (first statement of Th. 1). Then, we prove the regret bound with HLS representation (second statement of Th. 1).

**Regret bound for LINUCB without HLS condition.** Let  $\epsilon > 0$  to be chosen later. We start by splitting the regret into rounds where the minimum gap is above  $\epsilon$ ,

$$R_n = \underbrace{\sum_{t=1}^n \mathbb{1}\{\Delta(x_t) \geq \epsilon\} r_t}_{(a)} + \underbrace{\sum_{t=1}^n \mathbb{1}\{\Delta(x_t) < \epsilon\} r_t}_{(b)}. \quad (28)$$

From Lem. 12,

$$(a) \leq \frac{32\lambda(\Delta_{\max} S_\phi \sigma)^2}{\epsilon} (2 \log(1/\delta) + d_\phi \log(1 + nL_\phi^2 / (\lambda d_\phi)))^2. \quad (29)$$

Term (b) can be bounded using Lem. 14 with the sequence of events  $\{\Delta(x_t) < \epsilon\}_{t \geq 1}$ ,

$$(b) \leq \sqrt{8\Delta_{\max}^2 N_n d_\phi \log(1 + N_n L_\phi^2 / (\lambda d_\phi))} \left( \sqrt{\lambda} S_\phi + \sigma \sqrt{2 \log(1/\delta) + d_\phi \log(1 + nL_\phi^2 / (\lambda d_\phi))} \right). \quad (30)$$

It only remains to bound  $N_n$ , i.e., the count of these events. We have,

$$N_n := \sum_{t=1}^n \mathbb{1}\{\Delta(x_t) < \epsilon\} \leq \mathbb{P}(\Delta(x) < \epsilon) n + 2\sqrt{\mathbb{P}(\Delta(x) < \epsilon)(1 - \mathbb{P}(\Delta(x) < \epsilon)) n \log(2n/\delta)} + \frac{2}{3} \log(2n/\delta) \quad (31)$$

$$\leq C\epsilon^\alpha n + 2\sqrt{C\epsilon^\alpha(1 - C\epsilon^\alpha) n \log(2n/\delta)} + \frac{2}{3} \log(2n/\delta), \quad (32)$$

where the first inequality is from (any-time) Bernstein inequality and the second one is from the margin condition. The proof of the first result is concluded by setting  $\epsilon = n^{-1/\alpha}$ , for which the regret has the stated order. The final  $1 - 3\delta$  probability is the union bound of three events: the good event for LINUCB, concentration of contexts for Lem. 13, and concentration of “low-gap” contexts to bound  $N_n$ .

**Regret bound for LINUCB with HLS condition.** We follow similar steps as before, while considering a decreasing sequence  $\epsilon_t = t^{-1/\alpha}$  instead of a fixed  $\epsilon$ . We split the immediate regret at time  $t + 1$  as

$$r_{t+1} = \underbrace{\mathbb{1}\{\Delta(x_{t+1}) \geq \epsilon_{t+1}\} r_{t+1}}_{(a)} + \underbrace{\mathbb{1}\{\Delta(x_{t+1}) < \epsilon_{t+1}\} r_{t+1}}_{(b)}.$$

Let us start from (a). Since the minimum gap is larger than  $\epsilon_{t+1}$ , a sufficient condition for the immediate regret to be zero is (cf. proof of Lem. 10)

$$\frac{2L_\phi \beta_{t+1, \phi}(\delta)}{\sqrt{\lambda_{\min}(V_{t+1, \phi})}} < \epsilon_{t+1}. \quad (33)$$



Rearranging and expanding the definition of  $\epsilon_t$ ,

$$\lambda_{\min}(V_{t+1,\phi}) > 4L_\phi^2 \beta_{t+1,\phi}^2(\delta)(t+1)^{2/\alpha}. \quad (34)$$

Using Lem. 13 with  $\epsilon = \left(\frac{\lambda_{\phi,\text{HLS}}}{2L_\phi^2 C}\right)^{1/\alpha}$ ,

$$\lambda_{\min}(V_{t+1,\phi}) \geq \lambda + \frac{\lambda_{\phi,\text{HLS}}}{2}t - 6L_\phi^2 \sqrt{t \log(2d_\phi t/\delta)} - L_\phi^2 g_{t\phi}^\epsilon(\delta), \quad (35)$$

where  $g_{t\phi}^\epsilon(\delta) \leq O((C/\lambda_{\phi,\text{HLS}})^{2/\alpha}(\log t)^2)$  by our choice of  $\epsilon$ . Therefore, since (i)  $\alpha > 2$  and (ii) the minimum eigenvalue of  $V_{t+1,\phi}$  grows linearly with  $t$ , there exists a constant time  $\tau_\phi$  after which  $r_t = 0$ . Moreover, it is easy to see that  $\tau_\phi \propto (\lambda_{\phi,\text{HLS}})^{\frac{\alpha}{2-\alpha}}$ . Therefore, summing (a) from  $t = 1$  to  $t = n$  yields a total regret bounded by  $\Delta_{\max}\tau_\phi$ . It only remains to characterize the contribution of (b) to the regret. Using Lem. 14 on the sequence of events  $\{\Delta(x_t) < \epsilon_t\}$ ,

$$\sum_{t=1}^n \mathbb{1}\{\Delta(x_t) < \epsilon_t\} r_t \leq \sqrt{8\Delta_{\max}^2 N_n d_\phi \log(1 + N_n L_\phi^2 / (\lambda d_\phi))} \left( \sqrt{\lambda} S_\phi + \sigma \sqrt{2 \log(1/\delta) + d_\phi \log(1 + n L_\phi^2 / (\lambda d_\phi))} \right). \quad (36)$$

To bound  $N_n$ , note that  $\{\Delta(x_t) < \epsilon_t\}$  is a sequence of independent random variables. Moreover  $\mathbb{E}[N_n] = \sum_{t=1}^n \mathbb{P}\{\Delta(x_t) < \epsilon_t\} \leq C \sum_{t=1}^n \epsilon_t^\alpha$  and  $\mathbb{V}[N_n] = \sum_{t=1}^n \mathbb{P}\{\Delta(x_t) < \epsilon_t\}(1 - \mathbb{P}\{\Delta(x_t) < \epsilon_t\}) \leq C \sum_{t=1}^n \epsilon_t^\alpha$ . Thus, using Bernstein's inequality,

$$N_n \leq \mathbb{E}[N_n] + 2\sqrt{\mathbb{V}[N_n] \log(2n/\delta)} + \frac{2}{3} \log(2n/\delta) \leq C \sum_{t=1}^n \epsilon_t^\alpha + 2\sqrt{C \sum_{t=1}^n \epsilon_t^\alpha \log(2n/\delta)} + \frac{2}{3} \log(2n/\delta). \quad (37)$$

Since  $\sum_{t=1}^n \epsilon_t^\alpha = \sum_{t=1}^n 1/t \leq \log n + 1$ , we obtain the stated order of regret. The final  $1 - 3\delta$  probability is the union bound of the same three events considered in the proof of the first part of the theorem.

### C.5. HLS is Necessary and Sufficient

Prop. 2 shows that LINUCB suffers constant regret on any linear contextual bandit problem with strictly-positive minimum gap when the algorithm is run with a HLS realizable representation  $\phi$ . In this section, we prove that the HLS condition is also necessary for achieving constant regret. We do so by leveraging the asymptotic problem-dependent regret lower-bound for linear contextual bandits (Lattimore & Szepesvari, 2017; Hao et al., 2020; Tirinzoni et al., 2020), from which we show that any consistent algorithm (like LINUCB) must suffer logarithmic regret if the representation is not HLS. We note that Hao et al. (2020) proved that the asymptotic lower bound is sub-logarithmic when representation  $\phi$  is HLS, while we could not find a formal proof that the asymptotic regret is logarithmic when  $\phi$  is not HLS.

In addition to those in the main paper, we consider the same assumptions used to derive the asymptotic lower-bound for linear contextual bandits (Hao et al., 2020; Tirinzoni et al., 2020):

1. The set of contexts  $\mathcal{X}$  is finite;
2. The context distribution  $\rho$  is full-support, i.e.,  $\rho(x) > 0$  for each  $x \in \mathcal{X}$ ;
3. The reward noise  $\eta_t$  is i.i.d. from  $\mathcal{N}(0, 1)$ .

We note that the finiteness of  $\mathcal{X}$  implies that the minimum-gap assumption holds (i.e.,  $\Delta > 0$ ). We believe that the asymptotic lower-bound (and thus our results) could be generalized to infinite contexts with an assumption on the minimum positive gap, though this is out of the scope of this work.

We start by stating the asymptotic lower bound on the expected regret of any consistent bandit algorithm. Formally, we call an algorithm consistent if it suffers  $o(n^\alpha)$  regret for any  $\alpha \in (0, 1)$  in any linear contextual problem satisfying the assumptions above.

**Proposition 6** (Hao et al., 2020). *Consider any consistent bandit algorithm and any linear contextual bandit problem satisfying the assumptions above. Then,*

$$\liminf_{n \rightarrow \infty} \frac{\mathbb{E}[R_n]}{\log(n)} \geq v^*(\theta^*), \quad (38)$$

where  $v^*$  is the value of the optimization problem

$$\inf_{\eta(x,a) \geq 0} \sum_{x \in \mathcal{X}} \sum_{a \in \mathcal{A}} \eta(x,a) \Delta(x,a) \quad \text{s.t.} \quad \|\phi(x,a)\|_{V_\eta^{-1}}^2 \leq \frac{\Delta(x,a)^2}{2} \quad \forall x \in \mathcal{X}, a \in \mathcal{A} : \Delta(x,a) > 0, \quad (P)$$

where  $V_\eta := \sum_x \sum_a \eta(x,a) \phi(x,a) \phi(x,a)^T$ .

We now show that, if  $\phi$  is not HLS, then there exists a positive constant  $c > 0$  such that  $v^* \geq c$ . This implies that the asymptotic regret of any consistent strategy (including LINUCB) grows at rate at least  $c \log(n)$ , hence proving our main claim.

**Proposition 7.** *Consider any consistent bandit algorithm and any linear contextual bandit problem satisfying the assumptions above with a realizable representation  $\phi$  that is not HLS. Then, there exists a positive constant  $c > 0$  such that*

$$\liminf_{n \rightarrow \infty} \frac{\mathbb{E}[R_n]}{\log(n)} \geq c. \quad (39)$$

*Proof.* We start by noting that, if  $\phi$  is not HLS, then there exists at least one context  $\bar{x}$  and one sub-optimal arm  $\bar{a}$  such that  $\phi(\bar{x}, \bar{a}) \notin \text{span}(\{\phi^*(x)\}_{x \in \mathcal{X}})$ . Clearly, for any  $\lambda \geq 0$ ,  $v^*$  is larger than the value of the following optimization problem

$$\inf_{\eta(x,a) \geq 0} \sum_{x \in \mathcal{X}} \sum_{a \in \mathcal{A}} \eta(x,a) \Delta(x,a) \quad \text{s.t.} \quad \|\phi(\bar{x}, \bar{a})\|_{(\lambda I + V_\eta)^{-1}}^2 \leq \frac{\Delta(\bar{x}, \bar{a})^2}{2}. \quad (40)$$

This holds because we removed all constraints except the one for  $\bar{x}, \bar{a}$  and added a small regularization to the matrix  $V_\eta$  (which can only decrease the norm). Let  $\eta^*(x,a)$  be an optimal solution of this optimization problem. We now prove that there exists a positive constant  $c' > 0$  such that, for some context  $x$  and sub-optimal arm  $a$ ,  $\eta^*(x,a) \geq c'$ . Let us proceed by contradiction. Suppose that all sub-optimal arms have  $\eta^*(x,a) = 0$ . Since  $\eta^*$  is feasible, we must have that

$$\|\phi(\bar{x}, \bar{a})\|_{(\lambda I + \sum_{x \in \mathcal{X}} \eta(x, a^*(x)) \phi^*(x) \phi^*(x)^T)^{-1}}^2 \leq \frac{\Delta(\bar{x}, \bar{a})^2}{2}. \quad (41)$$

Let  $\lambda_i, u_i$  denote the eigenvalues/eigenvectors of the matrix  $\sum_{x \in \mathcal{X}} \eta(x, a^*(x)) \phi^*(x) \phi^*(x)^T$ . Note that at least one of the eigenvalues must be zero since the optimal features do not span  $\mathbb{R}^d$ . Moreover, since  $\phi(\bar{x}, \bar{a})$  is not in the span of the optimal arms, Lem. 28 ensures that there exists some  $\epsilon > 0$  such that  $|\phi(\bar{x}, \bar{a})^T u_i| \geq \epsilon$  for at least one eigenvector  $u_i$  associated with a zero eigenvalue. Then, the norm can be re-written as

$$\|\phi(\bar{x}, \bar{a})\|_{(\lambda I + \sum_{x \in \mathcal{X}} \eta(x, a^*(x)) \phi^*(x) \phi^*(x)^T)^{-1}}^2 = \sum_{i \in [d]} \frac{(\phi(\bar{x}, \bar{a})^T u_i)^2}{\lambda + \lambda_i} \geq \frac{\epsilon^2}{\lambda}. \quad (42)$$

Since  $\lambda$  was arbitrary, we can set it to any value  $\lambda < \frac{2\epsilon^2}{\Delta(\bar{x}, \bar{a})^2}$ , for which we have a contradiction. Thus, we have proved that there always exists a positive constant  $c' > 0$  such that, for some context  $x$  and sub-optimal arm  $a$ ,  $\eta^*(x,a) \geq c'$ . This means that the value of the original optimization problem is at least  $v^* \geq c' \Delta > 0$ . This concludes the proof.  $\square$

### C.6. An Explicit Bound on $\tau$ from Lemma 2

We need to find a  $t$  that satisfies:

$$t > \frac{1}{\lambda_{HLS}} \left( 8L^2 \sqrt{t \log(2dt/\delta)} + L^2 g_t(\delta) + \frac{4L^2 \beta_t(\delta)}{\Delta^2} - \lambda \right).$$

Note that:

$$\frac{4L^2 \beta_t(\delta)}{\Delta^2} < \frac{4L^2}{\Delta^2} \left( \lambda S^2 + \sigma^2 \log \left( \frac{d\lambda + tL^2}{\delta d} \right) \right) < L^2 g_t(\delta),$$

thus a sufficient condition is:

$$t > \frac{1}{\lambda_{HLS}} \left( \underbrace{8L^2 \sqrt{t \log(2dt/\delta)}}_{(A)} + \underbrace{2L^2 g_t(\delta)}_{(B)} \right). \quad (43)$$

We consider two cases.

**Case 1.** We first consider the case  $(A) \geq (B)$ . In this case, we just need to find a  $t$  such that:

$$t > \frac{1}{\lambda_{HLS}} 16L^2 \sqrt{t \log(2dt/\delta)} \quad \iff \quad t > \frac{256L^4}{\lambda_{HLS}^2} \log(2dt/\delta).$$

Inequalities of the form  $at \geq \log(bt)$  (with  $a, b > 0$ ) can be solved using the Lambert  $W$  function as follows:

$$\begin{aligned} e^{\ln(bt)-at} &\leq 1 \\ \iff bte^{-at} &\leq 1 \\ \iff -ate^{-at} &\geq -a/b := x. \end{aligned}$$

We get real solutions only if  $x > -e^{-1}$ . If  $x \in (e^{-1}, 0)$ , then  $-at = W_{-1}(x)$ . If  $x > 0$ , then  $-at = W_0(x)$ . We can make the bound more explicit by noting that  $-1 - \sqrt{2u} - u \leq W_{-1}(-e^{-u-1}) \leq -1 - \sqrt{2u} - 2u/3$  for  $u > 0$  (Chatzigeorgiou, 2016) and, when  $x > e$ ,  $\ln(x) - \ln \ln(x) \leq W_0(x) \leq \ln(x) - \frac{1}{2} \ln \ln(x)$  (Hoorfar & Hassani, 2008).

Since  $a = \frac{\lambda_{HLS}^2}{256L^4}$  and  $b = 2d/\delta$ , we have that  $0 > -a/b = \frac{-\delta\lambda_{HLS}^2}{512dL^4} > e^{-1}$  for reasonable values of the constants. Then  $t \geq \frac{-W_{-1}(-a/b)}{a}$ , or  $t \geq \frac{1+\sqrt{2u+u}}{a}$  with  $u = \ln(b/a) + 1$ . Then a valid  $t$  is:

$$t \geq \frac{768L^4}{\lambda_{HLS}^2} \ln \left( \frac{512dL^4}{\delta\lambda_{HLS}^2} \right).$$

**Case 2:** We now consider the case  $(A) < (B)$  in (43). We seek a  $t$  such that:

$$t > \frac{4L^2 g_t(\delta)}{\lambda_{HLS}},$$

but notice that:

$$L^2 g_t(\delta) < \frac{64L^2 \sigma^2 \lambda S^2}{\Delta^2} 4d^2 \log^2 \left( \frac{d\lambda + tL^2}{\delta d} \right),$$

so equivalently:

$$t \geq \frac{256L^2 \sigma^2 \lambda S^2}{\lambda_{HLS} \Delta^2} 4d^2 \log^2 \left( \frac{d\lambda + tL^2}{\delta d} \right).$$

We have an inequality of the kind  $\sqrt{ct} = \ln(bt)$ . Let  $y = \sqrt{t}$  and  $a = \sqrt{c}$ , then  $ay = 2 \ln(by)$ , then:

$$0 = \ln(by) - \frac{ay}{2} \iff -\frac{a}{2b} = \frac{-ay}{2} e^{-ay/2}.$$

We have that  $a = \sqrt{\frac{\lambda_{HLS} \Delta^2}{1024d^2 L^2 \sigma^2 \lambda S^2}}$ , and since  $\frac{d\lambda + tL^2}{\delta d} < \frac{2dL^2 t}{\delta}$ , then  $b = \frac{2dL^2}{\delta}$ . Note that  $0 > -\frac{a}{2b} = -\frac{\Delta \delta \sqrt{\lambda_{HLS}}}{64d^2 L^3 \sigma S \sqrt{\lambda}} > e^{-1}$  for reasonable values of the constants. Then  $-\frac{ay}{2} = W_{-1}(-\frac{a}{2b})$  and:

$$y \geq \frac{6}{a} \ln(2b/a) = \frac{384dLS\sigma\sqrt{\lambda}}{\sqrt{\lambda_{HLS}\Delta}} \ln \left( \frac{64d^2 L^3 \sigma S \sqrt{\lambda}}{\sqrt{\lambda_{HLS}\Delta\delta}} \right),$$

so a valid  $t$  is:

$$t \geq \left( \frac{384dLS\sigma\sqrt{\lambda}}{\lambda_{HLS}\Delta} \ln \left( \frac{64d^2 L^3 \sigma S \sqrt{\lambda}}{\sqrt{\lambda_{HLS}\Delta\delta}} \right) \right)^2.$$

In conclusion we have that:

$$t \geq \max \left\{ \frac{384^2 d^2 L^2 S^2 \sigma^2 \lambda}{\lambda_{\text{HLS}} \Delta^2} \ln^2 \left( \frac{64 d^2 L^3 \sigma S \sqrt{\lambda}}{\sqrt{\lambda_{\text{HLS}} \Delta \delta}} \right), \frac{768 L^4}{\lambda_{\text{HLS}}^2} \ln \left( \frac{512 d L^4}{\delta \lambda_{\text{HLS}}^2} \right) \right\}. \quad (44)$$

## D. HLS Representations and Best-Arm Identification

In this section, we show that the HLS condition also enables LINUCB to solve best-arm identification (BAI) problems in the linear contextual bandit setting. More precisely, we show that, given a HLS representation, LINUCB equipped with a generalized likelihood-ratio test (e.g., [Hao et al., 2020](#); [Tirinzoni et al., 2020](#)) stops after finite number of steps and returns the optimal arms of each context. In fact, thanks to the HLS condition, LINUCB collects sufficient information about  $\theta_\phi^*$  by pulling the optimal arms alone, hence enabling us to prove that the algorithm stops after a finite time.

For each parameter  $\theta \in \mathbb{R}^d$ , let  $a_x^*(\theta) := \operatorname{argmax}_{a \in \mathcal{A}} \phi(x, a)^T \theta$ . We assume that the true optimal arms  $a_x^*(\theta^*)$  are unique for each context  $x \in \mathcal{X}$ , with corresponding optimal feature vector  $\phi^*(x)$ . Moreover, we shall remove the subscripts  $\phi$  to simplify the notation.

Consider the following variant of the LINUCB algorithm. At each step  $t$ , before choosing the next arm to pull, we perform the following test:

$$\inf_{\theta \in \bar{\Theta}_t} \|\theta_t - \theta\|_{V_t} > \beta_t(\delta), \quad (45)$$

where

$$\bar{\Theta}_t := \left\{ \theta \in \mathbb{R}^d \mid \exists x \in \mathcal{X} : a_x^*(\theta) \neq a_x^*(\theta_t) \right\}. \quad (46)$$

is the set of alternative parameters with respect to  $\theta_t$ , i.e., those where the optimal arm of at least one context differs from the one of  $\theta_t$ . This is the standard generalized likelihood-ratio test adopted for best-arm identification (e.g., [Degenne et al., 2020](#)) and in works focusing on asymptotic optimality (e.g., [Hao et al., 2020](#); [Tirinzoni et al., 2020](#)). For finite contexts, it is known that the test can be re-written in the convenient form (Cf. (79) from [Tirinzoni et al. \(2020\)](#)):

$$\min_{x \in \mathcal{X}} \min_{a \neq a_x^*(\theta_t)} \frac{(\phi(x, a_x^*(\theta_t)) - \phi(x, a))^T \theta_t}{\|\phi(x, a_x^*(\theta_t)) - \phi(x, a)\|_{V_t^{-1}}} > \beta_t(\delta). \quad (47)$$

Then, if the test triggers, we simply return the current least-square estimate  $\theta_t$ . Otherwise, we keep running LINUCB in its original form.

Using the standard confidence set derived for LINUCB ([Abbasi-Yadkori et al., 2011](#)), it is easy to prove the following result.

**Lemma 15** (LINUCB for BAI is  $\delta$ -correct). *If (45) holds at time  $t$ , then, with probability at least  $1 - \delta$ , for all  $x \in \mathcal{X}$ ,  $a_x^*(\theta^*) = a_x^*(\theta_t)$ .*

*Proof.* By contradiction, suppose the statement does not hold. This means that, for some context  $x \in \mathcal{X}$ , the true optimal arm  $a_x^*(\theta^*)$  is sub-optimal for  $\theta_t$ . By definition, this implies that  $\theta^* \in \bar{\Theta}_t$ , so that,

$$\|\theta_t - \theta^*\|_{V_t} \geq \inf_{\theta \in \bar{\Theta}_t} \|\theta_t - \theta\|_{V_t} > \beta_t(\delta),$$

which holds with probability at most  $\delta$  since  $\theta^*$  is contained in the confidence ellipsoid with probability at least  $1 - \delta$ .  $\square$

The following result shows that LINUCB run with a HLS representation and the generalized likelihood-ratio test stops in finite time and retrieves the true optimal arms with high probability.

**Lemma 16.** *Let  $\phi$  be HLS with  $\lambda_{\text{HLS}} := \lambda_{\min}(\mathbb{E}[\phi^*(x)\phi^*(x)^T]) > 0$  and  $\delta \in (0, 1)$ . Let  $\tau \geq 1$  be such that, for all  $t \geq \tau - 1$ ,*

$$\lambda + t\lambda_{\text{HLS}} - 8L^2 \sqrt{t \log(2dt/\delta)} - L^2 g_t(\delta) \geq \frac{16L^2 \beta_{t+1}^2(\delta)}{\Delta^2}.$$

*Then, with probability  $1 - \delta$ , LINUCB for BAI stops in at most  $\tau$  steps and returns a parameter whose optimal arms match the true ones.*

*Proof.* The fact that the returned parameter is correct was already proved in the previous lemma, so we only focus on showing that LINUCB eventually stops. In order to achieve so, we lower bound the left-hand side in (47) by a function that grows lineally in time. Fix any time step  $t$  such that  $t + 1 \geq \tau$ ,  $x \in \mathcal{X}$ ,  $a \neq a_{\theta_{t+1}}^*(s)$ . Then, using Lem. 9,

$$\|\phi(x, a_x^*(\theta_{t+1})) - \phi(x, a)\|_{V_{t+1}^{-1}}^2 \leq \frac{4L^2}{\lambda_{\min}(V_{t+1})} \leq \frac{4L^2}{\lambda + t\lambda_{\text{HLS}} - 8L^2\sqrt{t\log(2dt/\delta)} - L^2g_t(\delta)}.$$

Similarly,

$$\begin{aligned} (\phi(x, a_x^*(\theta_{t+1})) - \phi(x, a))^T \theta_{t+1} &\geq \Delta(s, a) - 2\beta_{t+1}(\delta) \max_{a' \in [K]} \|\phi(x, a')\|_{V_t^{-1}} \geq \Delta - \frac{2L\beta_{t+1}(\delta)}{\sqrt{\lambda_{\min}(V_{t+1})}} \\ &\geq \Delta - \frac{2L\beta_{t+1}(\delta)}{\sqrt{\lambda + t\lambda_{\text{HLS}} - 8L^2\sqrt{t\log(2dt/\delta)} - L^2g_t(\delta)}} \geq \frac{\Delta}{2}, \end{aligned}$$

where we used Lem. 9 and the fact that  $t + 1 \geq \tau$ . Then, the left-hand side of (47) is at least,

$$\frac{\Delta}{4L} \sqrt{\lambda + t\lambda_{\text{HLS}} - 8L^2\sqrt{t\log(2dt/\delta)} - L^2g_t(\delta)} \geq \beta_{t+1}(\delta),$$

where the last inequality holds for  $t + 1 \geq \tau$ . This means that the test triggers for  $t + 1 \geq \tau$ , which concludes the proof.  $\square$

## E. Representation Selection

In this section, we study the properties of LEADER when a set of  $M$  *realizable* representations is provided. Denote by  $a_t$  the action selected by LEADER at time  $t$ . The the instantaneous regret at time  $t$  is

$$r_t = \mu^*(x_t) - \mu(x_t, a_t). \quad (48)$$

Note that since all the representations are realizable, we have that  $r_t = \langle \theta_i^*, \phi_i(x_t, a^*) - \phi_i(x_t, a_t) \rangle$  for any  $i \in [M]$ .<sup>12</sup> We will address misspecified representations in App. F.

### E.1. Leveraging Good Representations

The following lemma establishes a key property of LEADER. We will use it to leverage access to HLS representations, but the same argument could be used to leverage other nice properties of candidate representations, or even to combine them (see App. E.2).

**Lemma 17.** *Consider a contextual bandit problem with reward  $\mu$ , context distribution  $\rho$  and  $\Delta > 0$ . Let  $(\phi_i)$  be a set of  $M$  realizable linear representations such that  $\max_{x, a} \|\phi_i(x, a)\|_2 \leq L_i$  and  $\|\theta_i^*\|_i \leq S_i$ . Then, with probability  $1 - \delta$ , for all  $t \geq 1$ , the instantaneous regret of LEADER run with confidence parameter  $\delta$  is:*

$$r_t \leq \min_{i \in [M]} \left\{ 2\beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}} \right\}.$$

*Proof.* Notice that LINUCB is entirely off-policy, in the sense that the quality of its parameter estimates is not affected by the fact of observing feedback from actions selected by an external rule. Since LEADER updates the estimates of all representations at each step, we can affirm that it runs  $M$  parallel LINUCB instances. Notice that each LINUCB instance is run with confidence parameter  $\delta/M$ , where  $\delta$  is the global confidence parameter of LEADER. For this reason, the good event  $\mathcal{G}_i(\delta/M)$  defined in Prop. 3 holds with probability  $1 - \delta/M$  for any  $i \in [M]$ . From a union bound, the intersection of this  $M$  events:

$$\mathcal{G}(\delta) := \left\{ \forall i \in [M], \forall t \geq 1, \|\theta_{ti} - \theta_i^*\|_{V_{ti}} \leq \beta_{ti}(\delta/M) \right\}, \quad (49)$$

holds with probability at least  $1 - \delta$ . Recall that  $\mathcal{C}_{ti}(\delta) = \{\theta \in \mathbb{R}^{d_i} \mid \|\theta_{ti} - \theta\|_{V_{ti}} \leq \beta_{ti}(\delta)\}$ . So  $\theta_i^* \in \mathcal{C}_{ti}(\delta/M)$  for all  $i \in [M]$  under  $\mathcal{G}(\delta)$ .

<sup>12</sup>Recall that we abbreviate  $\phi_i$  as just  $i$  in subscripts.



We proceed to bound the instantaneous regret at any time  $t$  under  $\mathcal{G}(\delta)$ . Since all the representations are realizable, for all  $a \in [K]$ ,  $\mu(x_t, a) = \phi_j(x_t, a)^\top \theta_j^*$  for any  $j \in [M]$ . So it is also true that, for any  $i \in [M]$ :

$$\begin{aligned} \phi_i^*(x_t)^\top \theta_i^* &= \mu^*(x_t) = \max_{a \in [K]} \mu(x, a) = \max_{a \in [K]} \min_{j \in [M]} \phi_j(x_t, a)^\top \theta_j^* \\ &\leq \max_{a \in [K]} \min_{j \in [M]} \max_{\theta \in \mathcal{C}_{tj}(\delta/M)} \phi_j(x_t, a)^\top \theta \end{aligned} \quad (50)$$

$$= \min_{j \in [M]} \max_{\theta \in \mathcal{C}_{tj}(\delta/M)} \phi_j(x_t, a_t)^\top \theta, \quad (51)$$

$$\begin{aligned} &\leq \max_{\theta \in \mathcal{C}_{ti}(\delta/M)} \phi_i(x_t, a_t)^\top \theta \\ &\leq \phi_i(x_t, a_t)^\top \tilde{\theta}_{ti}, \end{aligned} \quad (52)$$

where (50) is from  $\mathcal{G}(\delta)$ , (51) is from the arm-selection rule of LEADER, and  $\tilde{\theta}_{ti} := \arg \max_{\theta \in \mathcal{C}_{ti}(\delta/M)} \phi_i(x_t, a_t)^\top \theta$ . Finally, for any  $i \in [M]$ :

$$\begin{aligned} r_t &= \phi_i^*(x_t)^\top \theta_i^* - \phi_i(x_t, a_t)^\top \theta_i^* \\ &\leq \phi_i(x_t, a_t)^\top (\tilde{\theta}_{ti} - \theta_i^*) \end{aligned} \quad (53)$$

$$\begin{aligned} &\leq \left\| \tilde{\theta}_{ti} - \theta_i^* \right\|_{V_{ti}^{-1}} \|\phi(x_t, a_t)\|_{V_{ti}^{-1}} \\ &\leq 2\beta_{ti}(\delta/M) \|\phi(x_t, a_t)\|_{V_{ti}^{-1}}, \end{aligned} \quad (54)$$

where (53) is from (52), and the last inequality is from  $\mathcal{G}(\delta)$  and  $\tilde{\theta}_{ti} \in \mathcal{C}_{ti}(\delta/M)$ . Since (54) holds for all  $i \in [M]$  with overall probability  $1 - \delta$ , we can take a minimum over representations to conclude the proof.  $\square$

**Remark.** Inequality (52) clarifies in which sense LEADER is optimistic: every representation  $i \in [M]$  can overestimate the optimal reward of  $x_t$  with the reward of the action  $a_t$  selected by LEADER, using an optimistic parameter  $\tilde{\theta}_{ti}$  from its own confidence ellipsoid  $\mathcal{C}_{ti}(\delta/M)$ . In this sense,  $a_t$  is optimistic according to all representations, even if each representation alone could prescribe a different optimistic action, since  $\tilde{\theta}_{ti}$  may differ from  $\arg \max_{\theta \in \mathcal{C}_{ti}(\delta/M)} \max_{a \in [K]} \phi_i(x_t, a)^\top \theta$ , the optimistic parameter of LINUCB. The fact that these prescriptions may all differ from  $a_t$  implies that all representations may concur to the final action selection.

**Proof of Theorem 2.** The regret upper bound for LEADER established by Thm. 2 is a simple consequence of Lemma 17 and the results established for the single-representation case in previous sections.

*Proof.* From Lemma 17, under the good event  $\mathcal{G}$  defined in (49):

$$R_n = \sum_{t=1}^n r_t \leq \sum_{t=1}^n \min_{i \in [M]} \beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}} \leq \min_{i \in [M]} \sum_{t=1}^n \beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}}. \quad (55)$$

We can then bound  $\beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}}$  individually for each representation  $i \in [M]$ . First of all, we can proceed as in the proof of Proposition 1. In this case, the regret of LEADER is bounded by the smallest of the regret bounds (according to Proposition 1) of each LINUCB instance if it was run alone with confidence parameter  $\delta/M$ . For HLS representations, we can proceed as in the proof of Lemma 2 instead. For these representations, the time  $\tau_i$  after which  $\beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}} = 0$  is the same defined in Lemma 2, with  $\delta$  replaced by  $\delta/M$ . The final statement unifies the two cases by defining  $\tau_i = \infty$  for representations that are not HLS. The overall probability is  $1 - 2\delta$ , from a union bound. The first  $1 - \delta$  is for  $\mathcal{G}$ . The additional  $1 - \delta$  is to apply Lemma 9, instantiated with  $\delta \leftarrow \delta/M$ , to each HLS representation.  $\square$

**Remark.** By using the same arguments used for the single-representation case, one can easily obtain an expected-regret version of Theorem 2 (cf. App. C.3) and relax the zero-gap assumption into a margin condition (cf. App. C.4).

## E.2. Mixing representations

To prove Theorem 3, we first need to establish variants of the preliminary results from App. C.1 that hold without the HLS condition.

**Lemma 18** (Generalization of Lemma 9 in Loewner ordering). *Under the same assumptions of Proposition 3, assuming the good event  $\mathcal{G}_i(\delta)$  holds, with probability  $1 - \delta$ , for all  $t \geq 1$ :*

$$V_{t+1,i} \succeq \lambda I_{d_i} + t \mathbb{E}_{x \sim \rho} [\phi_i^*(x) \phi_i^*(x)^\top] - \left( 8L_i^2 \sqrt{t \log(2d_i t / \delta)} + L_i^2 g_{ti}(\delta) \right) I_{d_i},$$

where  $\succeq$  denotes the Loewner ordering.

*Proof.* First, we can rewrite  $V_{t+1,i}$  as

$$\begin{aligned} V_{t+1,i} &= \lambda I_{d_i} + \sum_{k=1}^t \phi_i(x_k, a_k) \phi_i(x_k, a_k)^\top \\ &= \lambda I_{d_i} + \sum_{k=1}^t \mathbb{1} \{a_k = a_{s_k}^*\} \phi_i^*(x_k) \phi_i^*(x_k)^\top + \sum_{k=1}^t \mathbb{1} \{a_k \neq a_{s_k}^*\} \phi_i(x_k, a_k) \phi_i(x_k, a_k)^\top \\ &\succeq \lambda I_{d_i} + \sum_{k=1}^t \phi_i^*(x_k) \phi_i^*(x_k)^\top - \sum_{k=1}^t \mathbb{1} \{a_k \neq a_{s_k}^*\} \phi_i^*(x_k) \phi_i^*(x_k)^\top, \end{aligned}$$

where the last inequality is due to the positive-semidefiniteness of  $\sum_{k=1}^t \mathbb{1} \{a_k \neq a_{s_k}^*\} \phi_i(x_k, a_k) \phi_i(x_k, a_k)^\top$ . Note that, for two positive semidefinite matrices  $A, B$ , we have  $A - B \succeq A - \lambda_{\max}(B)I$ . To see this, note that a sufficient condition for  $A - B$  to be greater than  $A - \lambda_{\max}(B)I$  in Loewner ordering is  $\lambda_{\min}(A - B - A + \lambda_{\max}(B)I) \geq 0$ , which is clearly true. Therefore,

$$\begin{aligned} V_{t+1,i} &\succeq \lambda I_{d_i} + \sum_{k=1}^t \phi_i^*(x_k) \phi_i^*(x_k)^\top - \lambda_{\max} \left( \sum_{k=1}^t \mathbb{1} \{a_k \neq a_{s_k}^*\} \phi_i^*(x_k) \phi_i^*(x_k)^\top \right) I_{d_i} \\ &\succeq \lambda I_{d_i} + \sum_{k=1}^t \phi_i^*(x_k) \phi_i^*(x_k)^\top - L_i^2 g_{ti}(\delta) I_{d_i}, \end{aligned}$$

where  $g_{ti}$  is the upper bound on suboptimal pulls from Prop. 5, which is valid under  $\mathcal{G}_i$ . We can now apply the same reasoning to lower-bound the sum of outer products of optimal feature vectors:

$$\begin{aligned} \sum_{k=1}^t \phi_i^*(x_k) \phi_i^*(x_k)^\top &= \sum_{k=1}^t \phi_i^*(x_k) \phi_i^*(x_k)^\top \pm \sum_{k=1}^t \mathbb{E}_{s \sim \rho} [\phi_i^*(s) \phi_i^*(s)^\top] \\ &\succeq t \mathbb{E}_{s \sim \rho} [\phi_i^*(s) \phi_i^*(s)^\top] - \lambda_{\max} \left( \sum_{k=1}^t \underbrace{\mathbb{E}_{s \sim \rho} [\phi_i^*(s) \phi_i^*(s)^\top] - \phi_i^*(x_k) \phi_i^*(x_k)^\top}_{X_k} \right) I_{d_i}. \end{aligned}$$

Clearly  $\mathbb{E}[X_k] = 0$  since  $s \sim \rho$ . Also  $X_k^2 \preceq \lambda_{\max}(X_k^2)I_{d_i} \preceq \|X_k\|^2 I \preceq 4L_i^4 I$  since  $X_k$  is symmetric. Hence, from matrix Azuma (Prop. 27), w.p.  $1 - \delta'_t$ :

$$\lambda_{\max} \left( \sum_{k=1}^t X_k \right) \leq 4L_i^2 \sqrt{2t \log(d/\delta'_t)}.$$

The proof is completed by union bound over time with  $\delta'_t = \delta/(2t^2)$ .  $\square$

The proof of the following key lemma provides some insights on how LEADER is able to mix representations.

**Lemma 19.** *Make the same assumptions of Theorem 3 and assume the good event  $\mathcal{G}(\delta)$  defined in (49) holds. With probability at least  $1 - \delta$ , for each representation  $i \in [M]$ , there exists a constant  $\tilde{\tau}_i$  such that, for all  $t \geq \tilde{\tau}_i$ , whenever  $(x_{t+1}, a_{t+1}) \in Z_i$ , the instantaneous regret of LEADER is  $r_{t+1} = 0$ . Moreover, for all  $t \geq \max_{i \in [M]} \tilde{\tau}_i$ ,  $r_{t+1} = 0$ .*

*Proof.* Recall from Definition 1 that  $M_i = \mathbb{E}_{x \sim \rho}[\phi_i^*(x)\phi_i^*(x)^\top]$ , and  $Z_i = \{(x, a) \in \mathcal{X} \times \mathcal{A} \mid \phi_i(x, a) \in \text{Im}(M_i)\}$ , where  $\text{Im}(M_i)$  denotes the column space of  $M_i$ . We will fix a representation  $i \in [M]$ , assume  $(x_{t+1}, a_{t+1}) \in Z_i$ , and show that  $r_{t+1} = 0$  if  $t \geq \tilde{\tau}_i$ . Since the  $M$  representations together satisfy the mixed-HLS condition,  $\bigcup_{i \in [M]} Z_i$  covers  $\mathcal{X} \times \mathcal{A}$  and  $r_{t+1} = 0$  after the maximal  $\tilde{\tau}_i$ .

From Lemma 18 (instantiated with  $\delta \leftarrow \delta/M$  for each representation, from which the  $1 - \delta$  overall probability):

$$V_{t+1,i} \succeq \lambda I_{d_i} + tM_i - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right) I_{d_i} := B_{ti}. \quad (56)$$

Note that  $B_{ti}$  is invertible for all  $t \geq T_i$  for a sufficiently large constant  $T_i$ . To see this, interpret the RHS as a time-dependent affine transformation of  $M_i$ , and note that those eigenvalues of  $M_i$  that are initially zero keep decreasing with time, while those that are initially nonzero start increasing forever after a finite time. Without loss of generality, we assume that  $T_i \leq \tilde{\tau}_i$ . Otherwise, it suffices to replace  $\tilde{\tau}_i$  with  $\max\{\tilde{\tau}_i, T_i\}$ .

Let  $\mathbf{v} := \phi_i(x_{t+1}, a_{t+1}) / \|\phi_i(x_{t+1}, a_{t+1})\|$ . From the Loewner ordering and the invertibility:

$$\mathbf{v}^\top V_{t+1,i}^{-1} \mathbf{v} \leq \mathbf{v}^\top B_{ti}^{-1} \mathbf{v}. \quad (57)$$

Since by assumption  $(x_{t+1}, a_{t+1}) \in Z_i$ ,  $\mathbf{v} \in \text{Im}(M_i)$ , hence  $\mathbf{v}$  belongs to the span of  $k$  orthonormal eigenvectors of  $M_i$ , where  $k$  is the number of nonzero eigenvalues of  $M_i$ . Note that the orthonormal eigenvectors of  $B_{ti}$  coincide with the orthonormal eigenvectors of  $M_i$  since  $B_{ti}$  is an affine transformation of  $M_i$ . This means that  $\mathbf{v}$  belongs to the span of  $k$  eigenvectors of  $B_{ti}$ . Also note that, for  $t \geq T_i$ , the nonzero eigenvalues of  $M_i$  correspond to *positive* eigenvalues of  $B_{ti}$ . This means that  $\mathbf{v}$  belongs to the span of  $k$  eigenvectors of  $B_{ti}$  having positive corresponding eigenvalues. The smallest of these eigenvalues is:

$$\lambda + t\lambda_i^+ - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right), \quad (58)$$

where  $\lambda_i^+$  denotes the smallest nonzero eigenvalue of  $M_i$ , and all of these eigenvalues are upper-bounded by:

$$\lambda + tL_i^2 - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right). \quad (59)$$

Now from Lemma 29:

$$\sqrt{\mathbf{v}^\top V_{ti}^{-1} \mathbf{v}} \leq \sqrt{\mathbf{v}^\top B_{ti}^{-1} \mathbf{v}} \leq \frac{\lambda + tL_i^2 - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right)}{\lambda + t\lambda_i^+ - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right)} \frac{1}{\sqrt{\mathbf{v}^\top B_{ti} \mathbf{v}}}. \quad (60)$$

Finally, since  $\mathbf{v}$  is orthogonal to all the eigenvectors of  $B_{ti}$  that correspond to zero eigenvalues, from Lemma 30:

$$\mathbf{v}^\top B_{ti} \mathbf{v} \geq \lambda + t\mathbf{v}^\top M_i \mathbf{v} - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right) \quad (61)$$

$$\geq \lambda + t \min_{\substack{\mathbf{v} \in \text{Im}(M_i) \\ \|\mathbf{v}\|=1}} \mathbf{v}^\top M_i \mathbf{v} - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right) \quad (62)$$

$$\geq \lambda + t\lambda_i^+ - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right). \quad (63)$$

Finally, from Lemma 17, under the same good event  $\mathcal{G}(\delta)$ :

$$\begin{aligned} r_{t+1} &\leq 2\sqrt{\beta_{t+1,i}} \|\phi(x_{t+1}, a_{t+1})\|_{V_{t+1}^{-1}} \\ &\leq 2\sqrt{\beta_{t+1,i}} L_i \sqrt{\mathbf{v}^\top V_{ti}^{-1} \mathbf{v}} \\ &\leq \frac{2L_i \sqrt{\beta_{t+1,i}} \left[ \lambda + tL_i^2 - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right) \right]}{\left[ \lambda + t\lambda_i^+ - \left(8L_i^2 \sqrt{t \log(2Md_i t/\delta)} + L_i^2 g_{ti}(\delta/M)\right) \right]^{3/2}} \\ &= \frac{\tilde{O}(t)}{\Omega(t\sqrt{t}) - \tilde{O}(t)} = \tilde{O}\left(\frac{1}{\sqrt{t}}\right). \end{aligned} \quad (64)$$

Hence we can find  $\tilde{\tau}_i$  such that, for all  $t \geq \tilde{\tau}_i$ ,  $r_{t+1} < \Delta$ . By definition of  $\Delta$ ,  $r_{t+1} = 0$ .  $\square$

**Remark.** Although we do not provide an explicit value for  $\tilde{\tau}_i$  here, by examining (64) we can conclude that  $\sqrt{\tilde{\tau}_i} \propto L_i^3/(\Delta(\lambda_i^+)^{3/2}) \geq L_i^2/(\Delta\sqrt{\lambda_i^+})$  hence  $\tilde{\tau}_i \propto L_i^4/(\Delta^2\lambda_i^+)$ . So  $\lambda_i^+$  (which is always nonzero for non-degenerate representations) has effectively replaced  $\lambda_{i,\text{HLS}}$  for the context-action pairs belonging to  $Z_i$ . In this sense, we can say that  $\phi_i$  is always "locally HLS" w.r.t.  $Z_i$ .

### Proof of Theorem 3.

*Proof.* Assume the good event  $\mathcal{G}(\delta)$  from (49) holds. Let  $\tau = \max_{i \in [M]} \tilde{\tau}_i$ , where  $\tilde{\tau}_i$  is defined as in Lemma 19. From the same lemma, with probability  $1 - \delta$ ,  $r_{t+1} = 0$  if  $t \geq \tau$ . Hence, if  $n \geq \tau$ :

$$R_n = \sum_{t=1}^n r_t \leq \sum_{t=1}^{\tau} r_t.$$

So, in any case, from Lemma 17:

$$R_n \leq \sum_{t=1}^{\min\{\tau, n\}} r_t \leq 2 \sum_{t=1}^{\min\{\tau, n\}} \min_{i \in [M]} \beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}} \leq 2 \min_{i \in [M]} \sum_{t=1}^{\min\{\tau, n\}} \beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}}.$$

We can bound  $\beta_{ti}(\delta/M) \|\phi_i(x_t, a_t)\|_{V_{ti}^{-1}}$  individually for each representation as in Proposition 1, obtaining the desired statement. The overall probability is  $1 - 2\delta$ , from a union bound. The first  $1 - \delta$  is for  $\mathcal{G}$ . The additional  $1 - \delta$  was required to apply Lemma 19.  $\square$

A corollary of Theorem 3 provides a generalization of Lemma 2 to the case of redundant representations, which is stated by (Hao et al., 2020) without an explicit proof.

**Corollary 1.** Consider a contextual bandit problem with realizable linear representation  $\phi_i$  such that, for all  $x \in \mathcal{X}$  and  $a \in \mathcal{A}$ :

$$\phi_i(x, a) \in \text{span}\{\phi_i^*(x) \mid x \in \text{supp}(\rho)\}. \quad (65)$$

Assume  $\Delta > 0$ ,  $\max_{x,a} \|\phi_i(x, a)\|_2 \leq L_i$  and  $\|\theta_i^*\|_2 \leq S_i$ . Then, with probability at least  $1 - 2\delta$ , the regret of LINUCB after  $n \geq 1$  steps is at most

$$R_n \leq \frac{32\lambda\Delta_{\max}^2 S_\phi^2 \sigma^2}{\Delta} \left( 2 \ln \left( \frac{1}{\delta} \right) + d_\phi \ln \left( 1 + \frac{\tilde{\tau}_i L_\phi^2}{\lambda d_\phi} \right) \right)^2,$$

where  $\tilde{\tau}_i$  is defined as in the proof of Lemma 19.

*Proof.* We will show that  $\phi$  by itself satisfies the mixed-HLS condition. Then, this is just Theorem 3 in the special case  $M = 1$ .

Let  $M_i$  and  $Z_i$  be as in Definition 1. Note that  $\text{Im}(M_i)$  is precisely  $\text{span}\{\phi_i^*(x) \mid x \in \text{supp}(\rho)\}$ . Hence, condition (65) is saying that  $Z_i = \mathcal{X} \times \mathcal{A}$ . But this means that  $\phi_i$  by itself satisfies the mixed-HLS condition.  $\square$

Corollary 1 that, even if  $\phi_i$  is not HLS, but it spans the same subspace spanned by the whole set of (redundant) features, LINUCB (hence, also LEADER) can achieve constant regret.

Finally, we provide some examples of sets of representations that have the mixed-HLS property.

**Example 1.** The following are two equivalent (non-HLS) 2-dimensional representations, with  $\theta_1^* = \theta_2^* = [1, 1]$ , for a problem with 2 contexts, 2 arms, and uniform context distribution. Optimal features are underlined.

$$\begin{array}{cccc} \underline{\phi_1(x_1, a_1)} = [2, 0] & \phi_1(x_1, a_2) = [1, 0] & \underline{\phi_2(x_1, a_1)} = [2, 0] & \phi_2(x_1, a_2) = [0, 1] \\ \underline{\phi_1(x_2, a_1)} = [2, 0] & \phi_1(x_2, a_2) = [0, 1] & \underline{\phi_2(x_2, a_1)} = [2, 0] & \phi_2(x_2, a_2) = [1, 0]. \end{array}$$

The first representation is HLS restricted to context  $x_1$ , i.e.,  $Z_1 = \{(x_1, a_1), (x_1, a_2), (x_2, a_1)\}$ . The second is HLS restricted to context  $x_2$ , i.e.,  $Z_2 = \{(x_1, a_1), (x_2, a_1), (x_2, a_2)\}$ . Since  $Z_1 \cup Z_2 = \mathcal{X} \times \mathcal{A}$ ,  $\{\phi_1, \phi_2\}$  is mixed-HLS.

**Example 2.** The following are three equivalent (non-HLS) 2-dimensional representations, with  $\theta_1^* = \theta_2^* = \theta_3^* = [1, 1]$ , for a problem with 2 contexts, 3 arms, and uniform context distribution. Optimal features are underlined.

$$\begin{aligned} \underline{\phi_1(x_1, a_1)} = [2, 0] \quad \phi_1(x_1, a_2) = [1, 0] \quad \underline{\phi_1(x_1, a_3)} = [0, 1] \quad \underline{\phi_2(x_1, a_1)} = [0, 2] \quad \underline{\phi_2(x_1, a_2)} = [1, 0] \quad \phi_2(x_1, a_3) = [0, 1] \\ \underline{\phi_1(x_2, a_1)} = [2, 0] \quad \underline{\phi_1(x_2, a_2)} = [0, 1] \quad \phi_1(x_2, a_3) = [1, 0] \quad \underline{\phi_2(x_2, a_1)} = [0, 2] \quad \phi_2(x_2, a_2) = [0, 1] \quad \underline{\phi_2(x_2, a_3)} = [1, 0]. \end{aligned}$$

In this case  $Z_1 = \{(x_1, a_1), (x_1, a_2), (x_2, a_1), (x_2, a_3)\}$  and  $Z_2 = \{(x_1, a_1), (x_1, a_3), (x_2, a_1), (x_2, a_2)\}$ . Since  $Z_1 \cup Z_2 = \mathcal{X} \times \mathcal{A}$ ,  $\{\phi_1, \phi_2\}$  is mixed-HLS.

**Example 3.** Consider the following 2-dimensional representation, with  $\theta_1^* = [1, 1]$ , for a problem with 2 contexts, 2 arms, and uniform context distribution.

$$\begin{aligned} \underline{\phi_1(x_1, a_1)} = [2, 0] \quad \phi_1(x_1, a_2) = [1, 0] \\ \underline{\phi_1(x_2, a_1)} = [2, 0] \quad \phi_1(x_2, a_2) = [1, 0]. \end{aligned}$$

It is redundant since the features only span  $\mathbb{R}^1$ . For the same reason, it is not HLS. However, optimal features also span  $\mathbb{R}^1$ . So, according to Corollary 1, LINUCB and LEADER can achieve constant regret with access to this distribution. Indeed,  $Z_1 = \mathcal{X} \times \mathcal{A}$  and  $\phi_1$  by itself satisfies the mixed-HLS condition. Redundancy is still a problem: the regret will scale with  $d = 2$  even if all features lie in a one-dimensional subspace. Redundant representations are the subject of study of *sparse bandits* (e.g., Abbasi-Yadkori et al., 2012).

## F. Dealing with Misspecified Representations

In this section, we introduce a variant of LEADER, called E-LEADER, that handles misspecified representations. More precisely, we consider the realizable setting (Agarwal et al., 2012a) where we are given a set of  $M$  representations such that at least one is realizable. Then, we combine LEADER with a statistical test that, under certain conditions, allows eliminating misspecified representations. Finally, we analyze the regret of E-LEADER and show that, under a minimum misspecification assumption, the algorithm essentially retains the same guarantees of LEADER for the subset of realizable representations, while suffering only constant regret for eliminating the misspecified ones. While we use a similar test as the one designed by Agarwal et al. (2012a) based on the mean squared error, our setting and analysis involve several additional complications. In fact, Agarwal et al. (2012a) consider a finite set of representations, each being a direct mapping from context and actions to rewards. On the other hand, in our case this mapping is given only after pairing some features  $\phi$  (i.e., our concept of representation) with some parameter  $\theta$ . Since these parameters are unknown and lie in a continuous space, our setting basically involves infinite possible representations (according to the definition of Agarwal et al. (2012a)).

### F.1. Formal Setting and Assumptions

We assume to have a set of representations  $\{\phi_i\}_{i \in [M]}$  such that

$$\mu(x, a) = \phi_i(x, a)^T \theta_i^* + f_i(x, a) \quad \forall i \in [M], x \in \mathcal{X}, a \in [K]. \quad (66)$$

As for the realizable setting, we suppose that  $\|\theta_i^*\|_2 \in \mathcal{B}_i := \{\theta \in \mathbb{R}^{d_i} : \|\theta\|_2 \leq S_i\}$  and  $\|\phi_i(x, a)\|_2 \leq L_i$  for all  $i \in [M]$ . Moreover, we suppose that the non-linear component  $f_i$  (i.e., the misspecification) satisfies  $|f_i(x, a)| \leq \|f_i\|_\infty := \epsilon_i \leq \epsilon$ . The learner knows the upper bounds  $S_i, L_i$ , while  $f_i, \epsilon_i, \epsilon$  are all unknown.

The following are our two key assumptions.

**Assumption 2 (Realizability).** *There exists a subset  $\mathcal{I}^* \in [M]$  with  $|\mathcal{I}^*| \geq 1$  such that  $\epsilon_i = 0$  for each  $i \in \mathcal{I}^*$*

**Assumption 3 (Minimum misspecification).** *For each misspecified representation  $i \in [M] \setminus \mathcal{I}^*$ , there exists  $\gamma_i > 0$  such that*

$$\min_{\theta \in \mathcal{B}_i} \min_{\pi \in \Pi} \mathbb{E}_{x \sim \rho} \left[ \left( \phi_i(x, \pi(x))^T \theta - \mu(x, \pi(x)) \right)^2 \right] \geq \gamma_i, \quad (67)$$

where  $\Pi$  is the set of all mappings from  $\mathcal{X}$  to  $[K]$ .

Intuitively, Asm. 3 requires that, no matter what actions the algorithm takes, in expectation over contexts no linear predictor can perfectly approximate the reward function. A sufficient condition for this to hold is that there exists some contexts which occur with positive probability and for which all actions are misspecified. Finally, as a technical assumption, we suppose that  $|\mu(x, a)| \leq 1$  and  $|y_t| \leq 1$  almost surely. This simplifies the analysis and the notation, though it can be easily relaxed.



**Algorithm 2** E-LEADER

---

**Input:** representations  $(\phi_i)_{i \in [M]}$  with values  $(L_i, S_i)_{i \in [M]}$ , regularization factor  $\lambda \geq 1$ , confidence level  $\delta \in (0, 1)$ .  
 Initialize  $V_{1i} = \lambda I_{d_i}$ ,  $\theta_{1i} = 0_{d_i}$  for each  $i \in [M]$  and  $\Phi_1 = [M]$   
**for**  $t = 1, \dots$  **do**  
   Observe context  $x_t$   
   Pull action  $a_t \in \operatorname{argmax}_{a \in [K]} \min_{i \in \Phi_t} \{U_{ti}(x_t, a)\}$   
   Observe reward  $r_t$  and, for each  $i \in [M]$   
    $V_{t+1,i} = V_{ti} + \phi_i(x_t, a_t)\phi_i(x_t, a_t)^\top$   
    $\theta_{t+1,i} = V_{t+1,i}^{-1} \sum_{l=1}^t \phi_i(x_l, a_l)r_l$   
    $\tilde{\theta}_{t+1,i} = \operatorname{argmin}_{\theta \in \mathcal{B}_i} \sum_{l=1}^t (\phi(x_l, a_l)^T \theta - r_l)^2$   
    $\Phi_{t+1} := \left\{ i \in \Phi_t : E_t(\phi_i, \tilde{\theta}_{t+1,i}) \leq \min_{j \in [M]} \min_{\theta \in \mathcal{B}_j} (E_t(\phi_j, \theta) + \alpha_{t+1,i}) \right\}$   
**end for**

---

**Additional Notation.** Finally, introduce some additional terms which will be recurrent in this section.

- $E_t(\phi, \theta) := \frac{1}{t} \sum_{k=1}^t (\phi(x_k, a_k)^T \theta - y_k)^2$ : mean squared error (MSE) of representation  $(\phi, \theta)$  at time  $t$ ;
- $\tilde{\theta}_{ti} := \operatorname{argmin}_{\theta \in \mathcal{B}_i} E_{t-1}(\phi_i, \theta)$ : least-squares estimate projected onto the ball of valid parameters;
- $\mathbb{E}_k$  and  $\mathbb{V}_k$ : expectation and variance conditioned on the full history up to (and not including) round  $k$ .

## F.2. Algorithm

The pseudo-code of elimination-based LEADER (E-LEADER), which combines LEADER with a statistical test for eliminating misspecified representations, is provided in Alg. 2.

The algorithm maintains a set of active representations  $\Phi_t$  defined as  $\Phi_1 := [M]$  and, for  $t \geq 2$ ,

$$\Phi_t := \left\{ i \in \Phi_{t-1} : E_{t-1}(\phi_i, \tilde{\theta}_{ti}) \leq \min_{j \in [M]} \min_{\theta \in \mathcal{B}_j} (E_{t-1}(\phi_j, \theta) + \alpha_{tj}) \right\}, \quad (68)$$

where

$$\alpha_{tj} := \frac{20}{t-1} \log \frac{8M^2(12L_j S_j (t-1))^{d_j} (t-1)^3}{\delta} + \frac{1}{t-1}. \quad (69)$$

Note that E-LEADER needs to compute a projected (onto the ball of valid parameters) version of the least-squares estimates to define the set  $\Phi_t$ . While we need this to simplify the analysis, in practice one can use the standard regularized least-squares estimates  $\theta_{ti}$  instead of  $\tilde{\theta}_{ti}$ . Finally, E-LEADER proceeds exactly as LEADER, by selecting, at each time step, the action with the minimum UCB across all active representations. Computationally, the main difference between the two algorithms is that E-LEADER needs to maintain estimates of the mean squares error  $E_t$ . These estimates can be updated incrementally when using regularized least squares, which makes E-LEADER as efficient as LEADER.

## F.3. Analysis

### F.3.1. REALIZABLE REPRESENTATIONS ARE NEVER ELIMINATED

The first three lemmas show that any realizable representation  $i \in \mathcal{I}^*$  is never eliminated from  $\Phi_t$  with high probability.

**Lemma 20.** *Let  $\phi, \theta \in \mathbb{R}^d$  be such that  $|\phi(s, a)^T \theta| \leq 1$ . Then, for each  $i \in \mathcal{I}^*$ ,  $t \geq 1$ , and  $\delta \in (0, 1)$ ,*

$$\mathbb{P} \left( E_t(\phi_i, \theta_i^*) > E_t(\phi, \theta) + \frac{20}{t} \log \frac{4t}{\delta} \right) \leq \delta. \quad (70)$$

*Proof.* This result can be obtained by reproducing the first part of the proof of Lemma 4.1 in (Agarwal et al., 2012a). We report the main steps for completeness.

Define  $g(x, a) := \phi(x, a)^T \theta$ ,  $g^*(x, a) := \phi_i(x, a)^T \theta_i^*$ , and  $Z_k := (g(x_k, a_k) - y_k)^2 - (g^*(x_k, a_k) - y_k)^2$ . Note that  $\{Z_k\}$  is a martingale difference sequence with  $|Z_k| \leq 4$ . Then, using Freedman's inequality (Lem. 26), with probability at least  $1 - \delta$ ,

$$\sum_{k=1}^t \mathbb{E}_k[Z_k] - \sum_{k=1}^t Z_k \leq 2 \sqrt{\sum_{k=1}^t \mathbb{V}_k[Z_k] \log \frac{4t}{\delta}} + 16 \log \frac{4t}{\delta}.$$

Using Lemma 4.2 of (Agarwal et al., 2012a), we have that  $\mathbb{V}_k[Z_k] \leq 4\mathbb{E}_k[Z_k]$ . Then, after a simple manipulation, the inequality above yields  $-\sum_{k=1}^t Z_k \leq 20 \log \frac{4t}{\delta}$ . The proof is concluded by noting that  $\sum_{k=1}^t Z_k = t(E_t(\phi, \theta) - E_t(\phi_i, \theta_i^*))$ .  $\square$

**Lemma 21.** For each  $\delta \in (0, 1)$ ,

$$\mathbb{P} \left( \exists t \geq 1, i \in \mathcal{I}^*, j \in [M], \theta \in \mathcal{B}_j : E_t(\phi_i, \theta_i^*) > E_t(\phi_j, \theta) + \frac{20}{t} \log \frac{8M^2(12L_j S_j t)^{d_j} t^3}{\delta} + \frac{1}{t} \right) \leq \delta. \quad (71)$$

*Proof.* We shall use a covering argument for each representation  $j \in [M]$ . First note that, for any  $\xi > 0$ , there always exists a finite set  $\mathcal{C}_j \subset \mathbb{R}^{d_j}$  of size at most  $(3S_j/\xi)^{d_j}$  such that, for each  $\theta \in \mathcal{B}_j$ , there exists  $\theta' \in \mathcal{C}_j$  with  $\|\theta - \theta'\|_2 \leq \xi$  (see e.g. Lemma 20.1 of (Lattimore & Szepesvári, 2020)). Moreover, suppose that all vectors in  $\mathcal{C}_j$  have  $\ell_2$ -norm bounded by  $S_j$  (otherwise we can always remove vectors with large norm). Now take any two vectors  $\theta, \theta' \in \mathcal{B}_j$  with  $\|\theta - \theta'\|_2 \leq \xi$ . We have

$$\begin{aligned} E_t(\phi_j, \theta) &= \frac{1}{t} \sum_{k=1}^t (\phi_j(x_k, a_k)^T \theta \pm \phi_j(x_k, a_k)^T \theta' - y_k)^2 \\ &= \frac{1}{t} \sum_{k=1}^t (\phi_j(x_k, a_k)^T (\theta - \theta'))^2 + \frac{1}{t} \sum_{k=1}^t (\phi_j(x_k, a_k)^T \theta' - y_k)^2 + \frac{2}{t} \sum_{k=1}^t (\phi_j(x_k, a_k)^T (\theta - \theta')) (\phi_j(x_k, a_k)^T \theta' - y_k) \\ &\geq E_t(\phi_j, \theta') + \frac{2}{t} \sum_{k=1}^t (\phi_j(x_k, a_k)^T (\theta - \theta')) \underbrace{(\phi_j(x_k, a_k)^T \theta' - y_k)}_{|\cdot| \leq 2} \\ &\geq E_t(\phi_j, \theta') - \frac{4}{t} \sum_{k=1}^t \|\phi_j(x_k, a_k)\|_2 \|\theta - \theta'\|_2 \geq E_t(\phi_j, \theta') - 4L_j \xi. \end{aligned}$$

Using  $\xi = \frac{1}{4L_j t}$ ,

$$\begin{aligned} &\mathbb{P} \left( \exists t \geq 1, i \in \mathcal{I}^*, j \in [M], \theta \in \mathcal{B}_j : E_t(\phi_i, \theta_i^*) > E_t(\phi_j, \theta) + \frac{20}{t} \log \frac{4t^3}{\delta'_t} + \frac{1}{t} \right) \\ &\leq \sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}^*} \sum_{j \in [M]} \mathbb{P} \left( \exists \theta \in \mathcal{B}_j : E_t(\phi_i, \theta_i^*) > E_t(\phi_j, \theta) + \frac{20}{t} \log \frac{4t^3}{\delta'_t} + \frac{1}{t} \right) \\ &\leq \sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}^*} \sum_{j \in [M]} \mathbb{P} \left( \exists \theta' \in \mathcal{C}_j : E_t(\phi_i, \theta_i^*) > E_t(\phi_j, \theta') - \frac{1}{t} + \frac{20}{t} \log \frac{4t^3}{\delta'_t} + \frac{1}{t} \right) \\ &\leq \sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}^*} \sum_{j \in [M]} \sum_{\theta' \in \mathcal{C}_j} \mathbb{P} \left( E_t(\phi_i, \theta_i^*) > E_t(\phi_j, \theta') + \frac{20}{t} \log \frac{4t^3}{\delta'_t} \right) \leq \sum_{t=1}^{\infty} \sum_{i \in \mathcal{I}^*} \sum_{j \in [M]} \sum_{\theta' \in \mathcal{C}_j} \frac{\delta'_t}{t^2} \\ &\leq M^2 \sum_{t=1}^{\infty} \frac{\delta'_t}{t^2} (12L_j S_j t)^{d_j}. \end{aligned}$$

Here the first inequality is from the union bound, the second one follows by relating  $\theta$  with its closest vector in the cover as above, the third one is from another union bound, the fourth one uses Lemma 20, and the last one is from the maximum size of the cover. The result follows by setting  $\delta'_t = \frac{\delta}{2M^2(12L_j S_j t)^{d_j}}$ .  $\square$

**Lemma 22.** [Validity of representation set] With probability at least  $1 - \delta$ , for each  $i \in \mathcal{I}^*$  and time  $t \geq 1$ ,  $i \in \Phi_t$ .

*Proof.* Clearly  $i \in \Phi_1$  by definition. From Lemma 21, with probability at least  $1 - \delta$ , for any  $i \in \mathcal{I}^*$  and  $t \geq 1$ ,

$$E_t(\phi_i, \theta_i^*) \leq \min_{j \in [M]} \min_{\theta \in \mathcal{B}_j} \left( E_t(\phi_j, \theta) + \frac{20}{t} \log \frac{8M^2(12L_j S_j t)^{d_j} t^3}{\delta} + \frac{1}{t} \right).$$

By definition,  $\tilde{\theta}_{t+1,i}$  is the vector in  $\mathcal{B}_i$  minimizing  $E_t(\phi_i, \theta)$ . Therefore, since  $\theta_i^* \in \mathcal{B}_i$ ,

$$E_t(\phi_i, \tilde{\theta}_{t+1,i}) \leq E_t(\phi_i, \theta_i^*),$$

which concludes the proof.  $\square$

### F.3.2. MISSPECIFIED REPRESENTATIONS ARE EVENTUALLY ELIMINATED

Next, we show that, thanks to Asm. 3, all the misspecified representations are eliminated from  $\Phi_t$  at some point.

**Lemma 23.** Let  $Z_k(\phi, \theta) := (\phi(x_k, a_k)^T \theta - y_k)^2$ . Define  $v_t := \sum_{k=1}^t \mathbb{E}_k[\mathbb{V}_k[y_k | x_k, a_k]]$  and  $b_t(\phi, \theta) := \sum_{k=1}^t \mathbb{E}_k \left[ (\phi(x_k, a_k)^T \theta - \mu(x_k, a_k))^2 \right]$  as the sum of conditional variances and biases, respectively. Then, for any  $\delta \in (0, 1)$ ,

$$\mathbb{P} \left( \exists t \geq 1, i \in [M], \theta \in \mathcal{B}_i : \left| \sum_{k=1}^t Z_k(\phi_i, \theta) - b_t(\phi_i, \theta) - v_t \right| > 4\sqrt{t \log \frac{4Mt(12tS_i L_i)^{d_i}}{\delta}} + 2 \right) \leq \delta \quad (72)$$

*Proof.* We follow a covering argument analogous to the one used in Lem. 21. It is easy to see that, if two parameters  $\theta, \theta' \in \mathcal{B}_i$  are  $\xi$ -close in  $\ell_2$ -norm, then

$$\sum_{k=1}^t |Z_k(\phi_i, \theta) - Z_k(\phi_i, \theta')| \leq \sum_{k=1}^t \underbrace{|\phi_i(x_k, a_k)^T(\theta + \theta') - 2y_k|}_{\leq 4} \underbrace{|\phi_i(x_k, a_k)^T(\theta - \theta')|}_{\leq L_i \xi} \leq 4L_i \xi t.$$

For each time  $t \geq 1$  and  $i \in [M]$ , we build a  $\xi_{ti}$ -cover  $\mathcal{C}_{ti}$  of  $\mathcal{B}_i$ , where  $\xi_{ti} = 1/(4L_i t)$ . Therefore,

$$\begin{aligned} & \mathbb{P} \left( \exists t \geq 1, i \in [M], \theta \in \mathcal{B}_i : \left| \sum_{k=1}^t Z_k(\phi_i, \theta) - \sum_{k=1}^t \mathbb{E}_k[Z_k(\phi_i, \theta)] \right| > 4\sqrt{t \log(2t/\delta'_t)} + 2 \right) \\ & \leq \sum_{t \geq 1} \sum_{i \in [M]} \sum_{\theta \in \mathcal{C}_{ti}} \mathbb{P} \left( \left| \sum_{k=1}^t Z_k(\phi_i, \theta) - \sum_{k=1}^t \mathbb{E}_k[Z_k(\phi_i, \theta)] \right| > 4\sqrt{t \log(2t/\delta'_t)} + 2 - 8L_i \xi_{ti} t \right) \\ & \leq M \sum_{t \geq 1} \delta'_t (12tS_i L_i)^{d_i}, \end{aligned}$$

where the first inequality is from a union bound and a reduction to the cover  $\mathcal{C}_{ti}$  (for both terms inside the absolute value) while the second one uses Azuma's inequality (Lem. 25). The proof is concluded by setting  $\delta'_t = \frac{\delta}{2M(12tS_i L_i)^{d_i}}$  and noting that

$$\mathbb{E}_k[Z_k(\phi_i, \theta)] = \mathbb{E}_k \left[ \mathbb{V}_k[y_k | x_k, a_k] + (\phi_i(x_k, a_k)^T \theta - \mu(x_k, a_k))^2 \right].$$

$\square$

**Lemma 24 (Elimination).** Let  $i \in [M]$  be any misspecified representation. Then, with probability at least  $1 - 2\delta$ , we have  $i \notin \Phi_{t+1}$  for all  $t$  such that

$$t \geq \frac{1}{\gamma_i} \min_{j \in \mathcal{I}^*} \left\{ 5 + 4\sqrt{t \log \frac{4Mt(12tS_i L_i)^{d_i}}{\delta}} + 4\sqrt{t \log \frac{4Mt(12tS_j L_j)^{d_j}}{\delta}} + 20 \log \frac{8M^2(12L_j S_j t)^{d_j} t^3}{\delta} \right\}.$$

*Proof.* Suppose that the misspecified representation  $i \in [M]$  is active at time  $t + 1$  (i.e.,  $i \in \Phi_{t+1}$ ). Let  $j \in \mathcal{I}^*$  be any realizable representation. Then,

$$\begin{aligned} E_t(\phi_i, \tilde{\theta}_{t+1,i}) &\leq \min_{j \in \Phi_{t+1}} \min_{\theta \in \mathcal{B}_j} (E_t(\phi_j, \theta) + \alpha_{t+1,i}) \leq E_t(\phi_j, \theta_j^*) + \alpha_{t+1,j} \\ &\leq \frac{v_t}{t} + 4\sqrt{\frac{1}{t} \log \frac{4Mt(12tS_jL_j)^{d_j}}{\delta}} + \frac{2}{t} + \alpha_{t+1,j}, \end{aligned}$$

where the first inequality uses the definition of  $\Phi_{t+1}$ , the second one uses that  $j$  is active at any time, and the last one uses Lem. 23 together with  $b_t(\phi_j, \theta_j^*) = 0$  since  $j$  is realizable. Similarly, by another application of Lem. 23 to the left-hand side,

$$E_t(\phi_i, \tilde{\theta}_{t+1,i}) \geq \frac{b_t(\phi_i, \tilde{\theta}_{t+1,i}) + v_t}{t} - 4\sqrt{\frac{1}{t} \log \frac{4Mt(12tS_iL_i)^{d_i}}{\delta}} - \frac{2}{t}.$$

Combining these two and expanding the definition of  $\alpha_{t+1,j}$ ,

$$b_t(\phi_i, \tilde{\theta}_{t+1,i}) \leq 5 + 4\sqrt{t \log \frac{4Mt(12tS_iL_i)^{d_i}}{\delta}} + 4\sqrt{t \log \frac{4Mt(12tS_jL_j)^{d_j}}{\delta}} + 20 \log \frac{8M^2(12L_jS_jt)^{d_j}t^3}{\delta}.$$

Finally, by Asm. 3

$$\begin{aligned} b_t(\phi_i, \tilde{\theta}_{t+1,i}) &\geq \min_{\theta \in \mathcal{B}_i} b_t(\phi_i, \theta) = \min_{\theta \in \mathcal{B}_i} \mathbb{E}_{x \sim \rho} \left[ \sum_{k=1}^t (\phi_i(x, \pi_k(x))^T \theta - \mu(x, \pi_k(x)))^2 \right] \\ &\geq t \min_{\theta \in \mathcal{B}_i} \min_{\pi \in \Pi} \mathbb{E}_{x \sim \rho} \left[ (\phi_i(x, \pi(x))^T \theta - \mu(x, \pi(x)))^2 \right] \geq t\gamma_i. \end{aligned}$$

Relating the last two inequalities, where in the former we can take a minimum over all realizable representations, we obtain an inequality of the form  $t\gamma_i \leq o(t)$ , from which we can find  $t$  such that  $i$  is eliminated. The result holds with probability  $1 - 2\delta$  due to the combination of two union bounds: one for the validity of the representation set, and the other for using Lem. 23. □

### F.3.3. REGRET ANALYSIS

Given that all misspecified representations are eliminated at some point, the analysis follows straightforwardly from the one of LEADER. In particular, if we let  $\Gamma_i$  be the first (deterministic) time step such that the inequality in Lem. 24 holds, we can bound the immediate regret for each  $t \leq \Gamma_i$  with its maximum value (i.e., 2). Then, the total regret of E-LEADER is

$$R_n = \sum_{t=1}^{\lceil \Gamma \rceil} r_t + \sum_{t=\lceil \Gamma \rceil+1}^n r_t \leq 2\lceil \Gamma \rceil + \sum_{t=\lceil \Gamma \rceil+1}^n r_t,$$

where  $\Gamma := \max_{i \in [M] \setminus \mathcal{I}^*} \Gamma_i$ . Then, the second term above can be bounded using exactly the same proof as LEADER since only the realizable representations remain active after  $\lceil \Gamma \rceil$ . Hence, E-LEADER suffers only constant regret for eliminating misspecified representations. Moreover, if at least one realizable representation is HLS, or the mixing HLS condition holds, the second term above is also constant, and E-LEADER suffers constant regret.

## F.4. Experiments

We finally report some numerical simulations where we compare E-LEADER with other model selection baselines (the one that performed better) on a toy problem with misspecification and on feature representations extracted from real data.

**Toy problem.** For the toy problem, we modified the experiment with varying dimensions of App. G.1.2. In addition to the representations considered in such experiment, we added four misspecified features, one with half the dimensionality of the base representation  $\phi_{\text{orig}}$  (see App. G.1.2), one with one third of the same, one generated randomly with dimension 3, and one generated randomly with dimension 9.

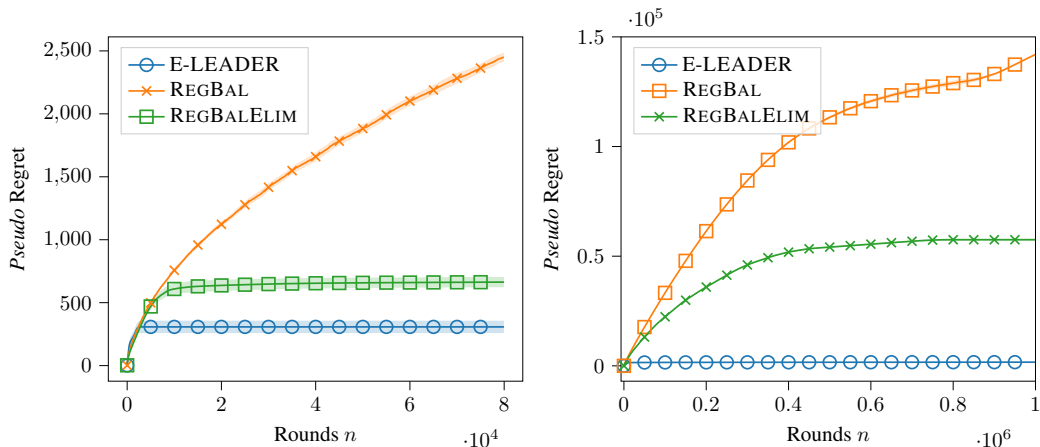


Figure 5. Regret of LEADER and model-selection baselines on two misspecified problems. (left) Toy problem with varying dimensions and misspecified representations. (right) 20 representations extracted from the Last.fm dataset.

**Last.fm dataset.** For the experiment with real data, we use the Last.fm dataset (Cantador et al., 2011), which is a list of users and music artists, together with information about the artists listened by each user. This dataset has been obtained from Last.fm (<https://www.last.fm/>) online music system. We first preprocessed the dataset by keeping only artists listened by at least 70 users and users that listened at least to 10 different artists. We thus obtained a dataset of 1322 users (which we treat as contexts) and 220 artists (which we treat as arms) reporting the number of times each user listened to each artist (which we treat as reward). Then, we generated multiple linear representations as follows. First, we extracted context-arm features via a low-rank factorization (taking only the highest 150 singular values) of the full matrix. Then, for each of 20 representations we generated a random neural network architecture by sampling the number of hidden layers in  $\{1, 2\}$ , the number of hidden neurons uniformly in  $[50, 200]$ , and the size of the output layer uniformly in  $[5, 50]$ . We obtained 20 neural networks with  $R^2$  score on a test set ranging from 0.5 to 0.85. For each of them, we took the last layer of the network as our linear model, thus obtaining 20 linear representations with varying dimension and misspecification. We found that 12 of these representations satisfy the HLS condition and took the best (in terms of  $R^2$ ) of these as our ground-truth model. Finally, we normalized all these representations to have  $\|\theta^*\|_2 = 1$  by properly re-scaling  $\phi$ .

**Results.** The results of running E-LEADER and the other model-selection baselines on the two problems described above are shown in Fig. 5. We compare to the baselines that perform best in the experiments of the main paper, i.e., REGBAL and REGBALELIM. For these baselines, we used  $d\sqrt{n}$  as the oracle upper bound to the  $n$ -step regret of LINUCB with a  $d$ -dimensional representation instead of the theoretical upper bound. Fig. 5 shows that E-LEADER outperforms the other baselines and quickly transitions to constant regret. REGBALELIM also transitions to constant regret since it eventually eliminates all misspecified representations but it takes an order of magnitude more steps than E-LEADER to achieve so.

## G. Experimental Details

In this section we provide further details on how the experiments from Section 6 were designed, the experimental setup, and additional results.

### G.1. Synthetic Problems with Finite Contexts

The two synthetic experiments reported in Section 6 and the motivating example from Figure 1 all pertain to the same randomly-generated contextual bandit problem with 20, 5 actions, uniform context distribution  $\rho$ , and noise standard deviation  $\sigma = 0.3$ . All the representations we consider are normalized so that  $S_i = \|\theta_i^*\| = 1$ , while  $L_i$  can change.

To construct the reward function  $\mu$  for our problem, we first generate a 6-dimensional random feature map  $\phi_{\text{orig}} : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^6$ . Specifically, each of the  $20 \times 5 \times 6$  elements of  $\phi_{\text{orig}}$  is sampled independently from a standard normal distribution. We separately sample a 6-dimensional parameter  $\theta_{\text{orig}}^*$ , each element from a uniform distribution over the interval  $[-1, 1]$ . We then normalize the parameter to have  $S_{\text{orig}} = 1$ . The reward function  $\mu$  is defined a posteriori as  $\mu(x, a) = \phi_{\text{orig}}(x, a)^\top \theta_{\text{orig}}^*$ ,



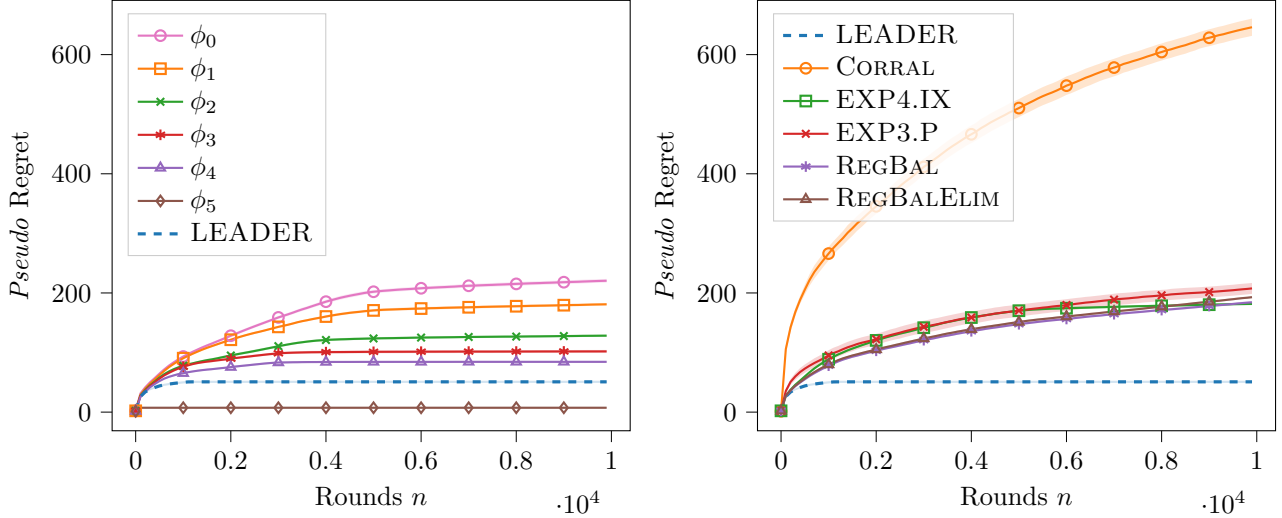


Figure 6. Regret of LEADER, LINUCB (left) with different representations and model-selection baselines (right) on the motivating representation-selection problem (App. G.1.1).

so  $\phi_{\text{orig}}$  is linearly realizable by construction.

From Lemma 8,  $\phi_{\text{orig}}^*$  is almost surely HLS. However, we verify this by checking that  $\lambda_{\text{orig,HLS}}$  is positive. All the representations we consider in the following are derived from and equivalent to  $\phi_{\text{orig}}$ .

### G.1.1. MOTIVATING EXAMPLE

The motivating example from Figure 1 is obtained from  $\phi_{\text{orig}}$  ( $\phi_5$  in the figure) by reducing the rank of  $\mathbb{E}_{x \sim \rho}[\phi_i^*(x)\phi_i^*(x)^\top]$  from 6 (full-rank), to 5, 4, 3, 2 and 1, using the procedure described in Lemma 7. A random linear transformation is then applied to each representation, which is also normalized to have  $S_i = 1$ . Notice that this operations do not affect the rank of  $\mathbb{E}_{x \sim \rho}[\phi_i^*(x)\phi_i^*(x)^\top]$  (Lemma 6). So we have a total of six representations, where only  $\phi_{\text{orig}}$  is HLS. The experiment is reproduced in Figure 6, where we also report the regret of the model-selection baselines. The results are in line with the theory. Only LINUCB with the HLS representation and LEADER achieve constant regret, significantly outperforming the other LINUCB instances.

### G.1.2. VARYING DIMENSION

Starting from  $\phi_{\text{orig}}$ , we construct five equivalent representations with dimension ranging from 2 to 6. To reduce the dimension, we just replace a subset of features of  $\phi_{\text{orig}}$  (that is, elements of  $\phi_{\text{orig}}(x, a)$  for all  $x, a$ ) with a single feature that is a linear combination weighted by the corresponding elements of  $\theta_{\text{orig}}^*$ . To obtain an equivalent representation, it is enough to set the corresponding element of the new  $\theta_i^*$  to 1. We then apply the procedure from Lemma 7 to obtain  $\text{rank}(\mathbb{E}_{x \sim \rho}[\phi_i^*(x)\phi_i^*(x)^\top]) = 1$ . Each representation is then passed through a random linear transformation and normalized to have  $S_i = 1$ . So, of the seven representations, only the original  $\phi_{\text{orig}}$  is HLS. The experiment is reproduced in Figure 7, where we also report the regret of the model-selection baselines.

### G.1.3. MIXING REPRESENTATIONS

For this experiment, we derive from  $\phi_{\text{orig}}$  six non-HLS 6-dimensional representations with the mixed-HLS property. The original HLS representation itself is *not* included in the set of candidates, to make mixing representations necessary.

To obtain the desired property, we remove two optimal features from each representation, replacing them with two copies of a new one that is an average of the two, similarly to what we did for the previous experiment, but without affecting the dimensionality. This operation makes the new representation non-HLS. However, since we only modify features of optimal arms, the new representation is still non-redundant. We pick the features to remove in a way that ensures the six representations together still have all the original features. Since the original  $\phi_{\text{orig}}$  was HLS, it is easy to verify that the set of representations constructed in this way has the mixed-HLS property from Definition 1. Again, each representation is

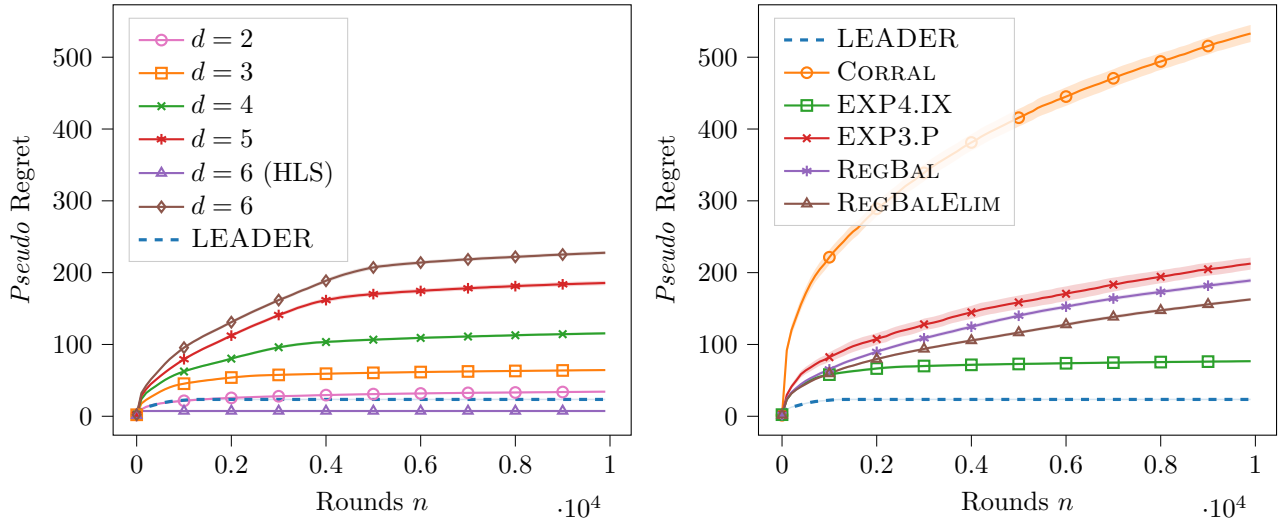


Figure 7. Regret of LEADER, LINUCB (left) with different representations and model-selection baselines (right) on the representation-selection problem with varying dimension (App. G.1.2).




Figure 8. Regret of LEADER, LINUCB (left) with different representations and model-selection baselines (right) on the problem of mixing representations (App. G.1.3).




Figure 9. Regret of LINUCB with two different representations on the synthetic problem with continuous contexts (App. G.2).

passed through a random linear transformation and normalized to have  $S_i = 1$ . The experiment is reproduced in Figure 8. Differently from the experiment reported in Figure 4 in the main paper, here *all* the base algorithms of REGBAL and REGBALELIM are updated at each time step for more fair comparison.

## G.2. Synthetic Problem with Continuous Contexts

In this section, we test LEADER on a synthetic contextual problem with continuous contexts. This example also clarifies the relationship between feature maps and context distributions in determining the goodness of a representation, showing how a change representation can sometimes correct for a ”bad” distribution.

The continuous context space is  $\mathcal{X} = \{x \in \mathbb{R}^d \mid \|x\| \leq 1\}$ . The finite action set is  $\mathcal{A} = \{[0, 0]^\top, [0, 1]^\top, [1, 0]^\top, [1, 1]^\top\}$ <sup>13</sup>. The context distribution  $\rho$  is uniform over  $\{(x_{[1]}, x_{[2]}) \in X \mid x_{[2]} \leq 0\}$ , where we use bracketed subscripts to index elements of vectors. Notice that  $\rho$  assigns zero probability to some contexts in  $\mathcal{X}$ . The reward function is:

$$\mu(x, a) = x_{[1]}a_{[1]} + x_{[2]}a_{[2]}.$$

Intuitively, each action selects one, both or none of the elements of the current context and receives the sum of the selected elements as a reward.

A natural 2-dimensional representation for this problem is:

$$\phi_1(x, a) = [x_{[1]}a_{[1]}, \quad x_{[2]}a_{[2]}]^\top,$$

with  $\theta_1^* = [1, 1]^\top$ . This would be a perfectly fine representation if  $\rho$  was full-support. However, since almost surely  $x_{[2]} \leq 0$ ,  $\phi_1^*(x)$  is either  $[0, 0]^\top$  or  $[x_{[1]}, 0]^\top$  depending on the sign of  $x_{[1]}$ . Hence, optimal features do not span  $\mathbb{R}^2$ , and  $\phi_1$  is not HLS.

Consider another representation for the same problem, this time 3-dimensional:

$$\phi_2(x, a) = [x_{[1]}a_{[1]} - x_{[1]}, \quad x_{[2]}a_{[2]} - x_{[2]}, \quad x_{[1]} + x_{[2]}]^\top,$$

with  $\theta_2^* = [1, 1, 1]^\top$ . It is easy to see that the two representations are equivalent. However, now  $\phi_2^*(x)$  can be  $[0, -x_{[2]}, x_{[1]} + x_{[2]}]^\top$  or  $[-x_{[1]}, -x_{[2]}, x_{[1]} + x_{[2]}]^\top$  depending on the sign of  $x_{[1]}$ , optimal features span  $\mathbb{R}^3$ , and  $\phi_2$  is HLS.

In Figure 9, we show the regret of LINUCB with the two representations ( $\lambda = 1$ ,  $\delta = 0.01$ ,  $\sigma = 0.2$ , 20 independent runs). Note that only  $\phi_2$  achieves constant regret.

<sup>13</sup>We directly assign a vector in  $\mathbb{R}^d$  to each one of the four actions for ease of notation.

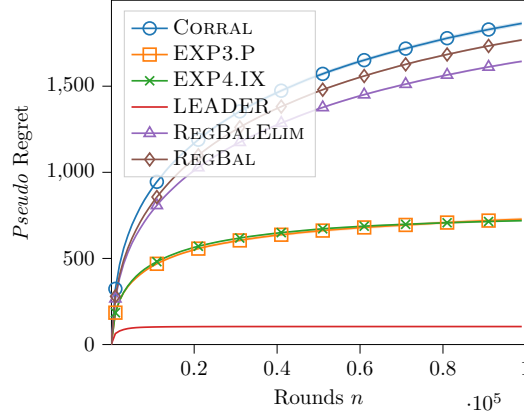


Figure 10. Regret of LEADER and model-selection baselines on 7 representations extracted from the Jester dataset.

### G.3. Jester Dataset

We report the comparison between LEADER and the model-selection baselines on the 7 feature representations extracted from the Jester dataset. For REGBAL and REGBALELIM, we use  $d\sqrt{n}$  as the oracle upper bound to the regret of LINUCB (instead of the theoretical one). Moreover, for all algorithms, we update all base learners with all collected samples (i.e., we share data across representations as LEADER does). The results are shown in Fig. 10. Consistently with the other experiments, LEADER outperforms all baselines and quickly transitions to constant regret thanks to the presence of multiple HLS representations.

### G.4. Hyperparameters for Model-Selection Baselines

In this section, we report hyperparameters and other settings used for the model-selection baselines. All base algorithms are always updated at each time step unless otherwise stated.

**EXP4.IX.** For the implicit-exploration parameter, we used the theoretical value  $\gamma = \sqrt{2 \log(M)/(nK)}$ . The learning rate is  $\eta = 2\gamma$ .

**CORRAL. and EXP3.P** We used the theoretical learning rate (respectively, exploration parameter) reported in App. A, with  $c(\delta) = \sqrt{d}(\sqrt{\lambda}S + \sigma\sqrt{d - 2 \log(\delta)})$ .

**Regret Balancing (with Elimination)** As a regret oracle, we used  $u : S_{t_i} \mapsto 4\beta_{t_i}(\delta)\sqrt{t \log(\det V_{t_i})}$  (for all bases).

## H. Auxiliary Results

**Lemma 25** (Azuma’s inequality). *Let  $\{(Z_t, \mathcal{F}_t)\}_{t \in \mathbb{N}}$  be a martingale difference sequence such that  $|Z_t| \leq a$  almost surely for all  $t \in \mathbb{N}$ . Then, for all  $\delta \in (0, 1)$ ,*

$$\mathbb{P} \left( \forall t \geq 1 : \left| \sum_{k=1}^t Z_k \right| \leq a\sqrt{t \log(2t/\delta)} \right) \geq 1 - \delta. \quad (73)$$

**Lemma 26** (Freedman’s inequality). *Let  $\{(Z_t, \mathcal{F}_t)\}_{t \in \mathbb{N}}$  be a martingale difference sequence such that  $|Z_t| \leq a$  almost surely for all  $t \in \mathbb{N}$ . Then, for all  $\delta \in (0, 1)$ ,*

$$\mathbb{P} \left( \forall t \geq 1 : \left| \sum_{k=1}^t Z_k \right| \leq 2\sqrt{\sum_{k=1}^t \mathbb{V}_k[Z_k] \log(4t/\delta)} + 4a \log(4t/\delta) \right) \geq 1 - \delta. \quad (74)$$

**Lemma 27** (Matrix Azuma, Tropp, 2012). *Let  $\{X_k\}_{k=1}^t$  be a finite adapted sequence of symmetric matrices of dimension  $d$ , and  $\{C_k\}_{k=1}^t$  a sequence of symmetric matrices such that for all  $k$ ,  $\mathbb{E}_k[X_k] = 0$  and  $X_k^2 \leq C_k^2$  almost surely. Then, with*

probability at least  $1 - \delta$ :

$$\lambda_{\max} \left( \sum_{k=1}^t X_k \right) \leq \sqrt{8\sigma^2 \log(d/\delta)}, \quad (75)$$

where  $\sigma^2 = \left\| \sum_{k=1}^t C_k^2 \right\|$ .

**Lemma 28.** *Let  $\Phi = [\phi_1, \dots, \phi_n]$ , where  $\phi_i \in \mathbb{R}^d$  for  $i \in [n]$ , be such that  $\text{span}(\Phi) \subset \mathbb{R}^d$ . Call  $\lambda_i$  the  $i$ -th eigenvalue (using an arbitrary order) of the matrix  $\Phi\Phi^T \in \mathbb{R}^{d \times d}$ , and let  $u_i$  be its corresponding eigenvector. Then, for any vector  $\phi \in \mathbb{R}^d$ , if  $\phi \notin \text{span}(\Phi)$ , there exists  $i \in [d]$  such that  $\lambda_i = 0$  and  $|\phi^T u_i| > 0$ .*

*Proof.* We prove the lemma by contradiction. Suppose that, for all  $i \in [d]$ , either  $\lambda_i > 0$  or  $|\phi^T u_i| = 0$ . Note that, since the columns of  $\Phi$  do not span  $\mathbb{R}^d$ ,  $\text{rank}(\Phi) = \text{rank}(\Phi\Phi^T) < d$ . This implies that at least one eigenvalue of  $\Phi\Phi^T$  is zero. Moreover, if  $\lambda_i = 0$ , the condition above implies that  $|\phi^T u_i| = 0$ . This means that  $\phi$  is orthogonal to all eigenvectors of  $\Phi$  associated to a zero eigenvalue. Define the matrix  $\bar{\Phi} := [\Phi, \phi] \in \mathbb{R}^{d \times (n+1)}$ . It must be that  $\text{rank}(\bar{\Phi}) > \text{rank}(\Phi)$  (otherwise  $\phi$  would lie in the span of the columns of  $\Phi$ ). Note that

$$\bar{\Phi}\bar{\Phi}^T = \sum_{i=1}^n \phi_i \phi_i^T + \phi \phi^T = \Phi\Phi^T + \phi \phi^T. \quad (76)$$

If  $u_i$  is an eigenvector of  $\Phi\Phi^T$  associated with a zero eigenvalue, then

$$\bar{\Phi}\bar{\Phi}^T u_i = \underbrace{\Phi\Phi^T u_i}_{=0} + \phi \underbrace{\phi^T u_i}_{=0} = 0 \quad (77)$$

since  $u_i$  is orthogonal to  $\phi$ . Hence,  $u_i$  is still an eigenvector of  $\bar{\Phi}\bar{\Phi}^T$  associated with eigenvalue zero. Therefore, the number of non-zero eigenvalues of  $\bar{\Phi}\bar{\Phi}^T$  is the same as the one of  $\Phi\Phi^T$ , which implies that  $\text{rank}(\bar{\Phi}) = \text{rank}(\bar{\Phi}\bar{\Phi}^T) = \text{rank}(\Phi\Phi^T) = \text{rank}(\Phi)$ . This yields the desired contradiction.  $\square$

**Lemma 29** (Kantorovich-like inequality). *Let  $\mathbf{v} \in \mathbb{R}^d$  with  $\|\mathbf{v}\| = 1$  and  $A \in \mathbb{R}^{d \times d}$  symmetric invertible with non-zero eigenvalues  $\lambda_1 \leq \dots \leq \lambda_d$  and corresponding orthonormal eigenvectors  $u_1, \dots, u_d$ . Let  $\mathcal{I} \subseteq [d]$  be any index set. If  $\mathbf{v} \in \text{span}\{u_i\}_{i \in \mathcal{I}}$  and  $\lambda_i > 0$  for all  $i \in \mathcal{I}$ :*

$$\mathbf{v}^T A^{-1} \mathbf{v} \leq \frac{(\max_{i \in \mathcal{I}} \lambda_i + \min_{i \in \mathcal{I}} \lambda_i)^2}{4 \max_{i \in \mathcal{I}} \lambda_i \min_{i \in \mathcal{I}} \lambda_i} \frac{1}{\mathbf{v}^T A \mathbf{v}}.$$

*Proof.* Since the eigenvectors are all orthogonal,  $\mathbf{v}^T u_i = 0$  for all  $i \notin \mathcal{I}$  (meaning  $i \in [d] \setminus \mathcal{I}$ ). Now consider the quadratic form:

$$\begin{aligned} \mathbf{v}^T A \mathbf{v} &= \mathbf{v}^T Q \Sigma Q^T \mathbf{v} && \text{(orthogonal diagonalization)} \\ &= \sum_{i=1}^d \lambda_i (\mathbf{v}^T u_i)^2 \\ &= \sum_{i \in \mathcal{I}} \lambda_i (\mathbf{v}^T u_i)^2 + \sum_{i \notin \mathcal{I}} \lambda_i (\mathbf{v}^T u_i)^2 && (78) \end{aligned}$$

$$\begin{aligned} &= \sum_{i \in \mathcal{I}} \lambda_i (\mathbf{v}^T u_i)^2 + \sum_{i \notin \mathcal{I}} \tilde{\lambda} (\mathbf{v}^T u_i)^2 && \text{(the second summation is zero anyway)} \\ &:= \mathbf{v}^T B \mathbf{v}, && (79) \end{aligned}$$

where  $\tilde{\lambda}$  is any eigenvalue from  $\{\lambda_i\}_{i \in \mathcal{I}}$  and  $B$  is the symmetric matrix of which (79) is the eigendecomposition. The substitution is legit because the dot products in the second summation in (78) are all zero. The new matrix  $B$  gives the same quadratic form but only has eigenvalues from  $\{\lambda_i\}_{i \in \mathcal{I}}$ . Since by hypothesis the surviving eigenvalues are all positive,  $B$  is

positive definite. We can do exactly the same for  $\mathbf{v}^\top A^{-1} \mathbf{v}$  since the orthonormal eigenvectors of  $A^{-1}$  are the same as  $A$ . Hence:

$$\begin{aligned} \mathbf{v}^\top A^{-1} \mathbf{v} \mathbf{v}^\top A \mathbf{v} &= \mathbf{v}^\top B^{-1} \mathbf{v} \mathbf{v}^\top B \mathbf{v} \\ &\leq \frac{(\lambda_{\max}(B) + \lambda_{\min}(B))^2}{4\lambda_{\max}(B)\lambda_{\min}(B)} \\ &= \frac{(\max_{i \in \mathcal{I}} \lambda_i + \min_{i \in \mathcal{I}} \lambda_i)^2}{4 \max_{i \in \mathcal{I}} \lambda_i \min_{i \in \mathcal{I}} \lambda_i}, \end{aligned} \tag{80}$$

where (80) is the standard Kantorovich matrix inequality (e.g., [Chen, 2013](#)). □

**Lemma 30.** *The smallest nonzero eigenvalue of symmetric p.s.d. matrix  $A \in \mathbb{R}^{d \times d}$  is:*

$$\min_{\substack{\mathbf{v} \in \text{Im}(A) \\ \|\mathbf{v}\|=1}} \mathbf{v}^\top A \mathbf{v},$$

where  $\text{Im}(A)$  denotes the column space of  $A$ .

*Proof.* Let  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_d$  be the eigenvalues of  $A$  with corresponding orthonormal eigenvectors  $u_1, u_2, \dots, u_d$ . The eigenvalues can be computed iteratively as:

$$\lambda_i = \min_{\substack{\mathbf{v} \in \{u_1, \dots, u_{i-1}\}^\perp \\ \|\mathbf{v}\|=1}} \mathbf{v}^\top A \mathbf{v}. \tag{81}$$

When  $\lambda_i$  is the smallest nonzero eigenvalue of  $A$ ,  $\text{span}\{u_1, \dots, u_{i-1}\}$  is precisely the solution set of:

$$A \mathbf{v} = \mathbf{0} \mathbf{v} = \mathbf{0}, \tag{82}$$

which is  $\ker(A)$ . Hence  $\{u_1, \dots, u_{i-1}\}^\perp = \ker(A)^\perp = \text{Im}(A)$  since  $A$  is symmetric. □