
Sparsity-Agnostic Linear Bandits with Adaptive Adversaries

Tianyuan Jin

Department of Electrical and Computer Engineering
National University of Singapore
Singapore
tianyuan@nus.edu.sg

Kyoungseok Jang

Dipartimento di Informatica
Università degli Studi di Milano
Milano, Italy
ksajks@gmail.com

Nicolò Cesa-Bianchi

Università degli Studi di Milano
Politecnico di Milano
Milano, Italy
nicolo.cesa-bianchi@unimi.it

Abstract

We study stochastic linear bandits where, in each round, the learner receives a set of actions (i.e., feature vectors), from which it chooses an element and obtains a stochastic reward. The expected reward is a fixed but unknown linear function of the chosen action. We study *sparse* regret bounds, that depend on the number S of non-zero coefficients in the linear reward function. Previous works focused on the case where S is known, or the action sets satisfy additional assumptions. In this work, we obtain the first sparse regret bounds that hold when S is unknown and the action sets are adversarially generated. Our techniques combine online to confidence set conversions with a novel randomized model selection approach over a hierarchy of nested confidence sets. When S is known, our analysis recovers state-of-the-art bounds for adversarial action sets. We also show that a variant of our approach, using Exp3 to dynamically select the confidence sets, can be used to improve the empirical performance of stochastic linear bandits while enjoying a regret bound with optimal dependence on the time horizon.

1 Introduction

K -armed bandits are a basic model of sequential decision-making in which a learner sequentially chooses which arm to pull in a set of K arms. After each pull, the learner only observes the reward returned by the chosen arm. After T pulls, the learner must obtain a total reward as close as possible to the reward obtained by always pulling the overall best arm. Linear bandits extend K -armed bandits to a setting in which arms belong to a d -dimensional feature space. In each round t of a linear bandit problem, the learner receives an action set $\mathcal{A}_t \subset \mathbb{R}^d$ from the environment, chooses an arm $A_t \in \mathcal{A}_t$ based on the past observations, and then receives a reward X_t . In this work, we consider the stochastic setting in which rewards are defined by $X_t = \langle \theta_*, A_t \rangle + \varepsilon_t$, where $\theta_* \in \mathbb{R}^d$ is a fixed latent parameter and ε_t is zero-mean independent noise. In linear bandits, the learner's goal is to minimize the difference between the total reward obtained by pulling in each round t the arm $a \in \mathcal{A}_t$ maximizing $\langle \theta_*, a \rangle$ and the total reward obtained by the learner.

In stochastic linear bandits, the regret after T rounds is known to be of order $d\sqrt{T}$ up to logarithmic factors. The linear dependence on the number d of features implies that the learner is better off by ignoring features corresponding to negligible components of the latent target vector θ_* . Hence,

one would like to design algorithms that depend on the number $S \ll d$ of relevant features without requiring any preliminary knowledge on θ_* . This is captured by the setting of sparse linear bandits, where θ_* is assumed to have only $0 < S \leq d$ nonzero components.

In the sparse setting, Lattimore and Szepesvári [17, Section 24.3] show that a regret of $\Omega(\sqrt{SdT})$ is unavoidable for any algorithm, even with knowledge of S . When S is known, this lower bound is matched (up to log factors) by an algorithm of Abbasi-Yadkori et al. [2] who, under the same assumptions and for the same algorithm, also prove an instance-dependent regret bound of $\tilde{O}(\frac{Sd}{\Delta})$. Here Δ is the minimum gap, over all T rounds, between the expected reward of the optimal arm and that of any suboptimal arm. In this work we focus on the sparsity-agnostic setting, i.e., when S is unknown. Fewer results are known for this case, and all of them rely on additional assumptions on the action set, or assumptions on the sparsity structure. For example, if the action set is stochastically generated, Oh et al. [22] prove a $\tilde{O}(S\sqrt{T})$ sparsity-agnostic regret bound. More recently, Dai et al. [9] showed a sparsity-agnostic bound $\tilde{O}(S^2\sqrt{T} + S\sqrt{dT})$ when the action set is fixed and equal to the unit sphere. In a model selection setting, Cutkosky et al. [8] prove a $\tilde{O}(S^2\sqrt{T})$ sparsity-agnostic regret bound for adversarial action sets, but under an additional nestedness assumption: $(\theta_*)_i \neq 0$ for $i = 1, \dots, S$. Surprisingly, no bounds improving on the $\tilde{O}(d\sqrt{T})$ regret of the OFUL algorithm [1] in the sparsity-agnostic case are known that avoid additional assumptions on the sparsity structure or on the action set generation.

Main contributions. Here is the summary of our main contributions. All the proofs of our results can be found in the appendix.

- We introduce a randomized sparsity-agnostic linear bandit algorithm, **SparseLinUCB**, achieving regret $\tilde{O}(S\sqrt{dT})$ with no assumptions on the sparsity structure (e.g., nestedness) or on the action set (which may be controlled by an adaptive adversary). When $S = o(\sqrt{d})$, our bound is strictly better than the OFUL bound $\tilde{O}(d\sqrt{T})$.
- Our analysis of **SparseLinUCB** simultaneously guarantees an instance-dependent regret bound $\tilde{O}(\max\{d^2, S^2d\}/\Delta)$, where Δ is the smallest suboptimality gap over the T rounds.
- If the sparsity level is known, our algorithm recovers the optimal bound $\tilde{\Theta}(\sqrt{SdT})$.
- We also introduce **AdaLinUCB**, a variant of **SparseLinUCB** that uses Exp3 to learn the probability distribution over a hierarchy of confidence sets in stochastic linear bandits. Unlike previous works, which only showed a $\tilde{O}(T^{2/3})$ regret bound for similar approaches, **AdaLinUCB** has a $\tilde{O}(\sqrt{T})$ regret bound. In experiments on synthetic data, **AdaLinUCB** performs better than **OFUL**.

Technical challenges. Recall that the arm chosen in each round by **OFUL** is

$$A_t = \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a, \hat{\theta}_t \rangle + \sqrt{\gamma_t} \|a\|_{V_{t-1}^{-1}} \quad (1.1)$$

where $\hat{\theta}_t$ is the regularized least-squares estimate of θ_* , $V_{t-1} = I + \sum_{s < t} A_s A_s^\top$ is the regularized covariance matrix of past actions, and $\sqrt{\gamma_t}$ is the radius of the confidence set

$$\left\{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_{t-1}\|_{V_{t-1}}^2 \leq \gamma_t \right\}. \quad (1.2)$$

The squared radius $\gamma_t = O(d \ln t)$ is such that θ_* belongs to (1.2) with high probability simultaneously for all $t \geq 1$. Our approach, instead, builds on the online to confidence set conversion technique of Abbasi-Yadkori et al. [2], where they show how to design a different confidence set for θ_* based on the predictions of an arbitrary algorithm for online linear regression, such that the squared radius of the confidence set is roughly equal to the regret bound of the algorithm. Using the algorithm **SeqSEW** for sparse online linear regression [13], whose regret bound is $O(S \log T)$, they obtain the optimal regret $\tilde{O}(\sqrt{SdT})$ for sparse linear bandits. Unfortunately, this result requires knowing S to properly set the radius of the confidence set. Our strategy **SparseLinUCB** (Algorithm 1) bypasses this problem by running the online to confidence set conversion technique over a hierarchy of nested confidence sets with radii $\alpha_i = 2^i \log T$ for $i = 1, \dots, n = \Theta(\log d)$. The framework of Abbasi-Yadkori et al. [2] guarantees that, for any sparsity value $S \in [d]$, there is a critical radius $\alpha_o = O(S \log T)$ such that, with high probability, θ_* lies in the set with radius α_i for all $i \geq o$. **SparseLinUCB** randomizes the choice of the index i of the confidence radius α_i , used for selecting the action at time t . If

Table 1: Comparison with other sparse linear bandit works. $S \in [d]$ is the sparsity level and Δ is the suboptimality gap (3.3). The nested assumption refers to $(\theta_*)_i \neq 0$ for $i = 1, \dots, S$. The minimum signal and the compatibility condition refer to assumptions on the distribution of the action set and on the smallest value of the non-zero elements in θ_* . Smoothed adversary refers to adversarially selected action sets with added Gaussian noise. The regret bounds listed in [1, 2, 23, 8, 25, 18, 9] are high-probability bounds: with high probability, the regret is of the same order as the bound in the table.

Reference	Sparsity Agnostic	Adaptive Adversary	Expected Regret	Assumptions
Abbasi-Yadkori et al. [1]	✓	✓	$\tilde{O}(\min\{d\sqrt{T}, d^2/\Delta\})$	-
Abbasi-Yadkori et al. [2]	✗	✓	$\tilde{O}(\min\{\sqrt{SdT}, dS/\Delta\})$	-
Pacchiano et al. [23, 24]	✓	✗	$\tilde{O}(S^2\sqrt{T})$	Nested, i.i.d. actions
Cutkosky et al. [8]	✓	✓	$\tilde{O}(S^2\sqrt{T})$	Nested
Pacchiano et al. [25]	✓	✓	$\tilde{O}(S^2\sqrt{T})$	Nested
	✓	✗	$\tilde{O}(S^2d^2/\Delta)$	Nested, i.i.d. actions
				Action set is hypercube
Lattimore et al. [18]	✗	✗	$\tilde{O}(S\sqrt{T})$	$\varepsilon_t \in [-1, 1]$
Sivakumar et al. [26]	✗	✓	$\tilde{O}(S\sqrt{T})$	Smoothed adversary
				Actions set spans \mathbb{R}^d
Hao et al. [15]	✗	✗	$\tilde{O}(\sqrt{ST})$	Minimum signal
Oh et al. [22]	✓	✗	$\tilde{O}(S\sqrt{T})$	Compatibility
Dai et al. [9]	✓	✗	$\tilde{O}(S^2\sqrt{T} + S\sqrt{dT})$	Action set is unit sphere
Lower bound [17]	✗	✓	$\Omega(\sqrt{SdT})$	-
This paper	✓	✓	$\tilde{O}\left(\min\{S\sqrt{dT}, \frac{1}{\Delta} \max\{d^2, S^2d\}\}\right)$	-
This paper	✗	✓	$\tilde{O}(\sqrt{SdT})$	-

the random index I_t is such that $I_t \geq o$, then we can bound the regret incurred at step t using standard techniques [1, 2]. By choosing $\mathbb{P}(I_t = i)$ proportional to 2^{-i} , we make sure that larger confidence sets (delivering suboptimal regret bounds) are chosen with exponentially small probability. If $I_t < o$, then θ_* is not guaranteed to lie in the confidence set of radius α_{I_t} with high probability. Our main technical contribution is to show that the regret summed over these bad rounds is bounded by $\tilde{O}(\sqrt{SdT}/Q)$, where $Q = \mathbb{P}(I_t \geq o)$. The proof of this bound requires showing that the regret in a bad round t (when $I_t < o$) can be bounded by $\sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}}$. Proving that $\|A_t^o\|_{V_{t-1}^{-1}}$ shrinks fast enough uses the fact that $\det V_t$ grows fast enough due to the exploration in the good rounds t (when $I_t \geq o$). This is done by a carefully designed peeling technique that partitions $[T]$ in blocks based on the value of $\det V_t$.

To extend our analysis of SparseLinUCB and obtain instance-dependent regret bound, we apply the techniques of Abbasi-Yadkori et al. [1] to show that the regret over the good rounds is bounded by $\tilde{O}(d^2/\Delta)$. The regret over a bad round t is controlled by $(\alpha_o/\Delta) \|A_t^o\|_{V_{t-1}^{-1}}^2$ and—using techniques similar to the instance-independent analysis—we bound the regret summed over all bad rounds with $\tilde{O}(Sd/(Q\Delta))$.

Given that SparseLinUCB uses a fixed probability of order 2^{-i} to choose its confidence radius α_i , it is tempting to explore adaptive probability assignments, that increase the probability of a confidence set proportionally to the rewards obtained by the actions that were selected based on that set. Algorithm AdaLinUCB (see Algorithm 2) is a variant of SparseLinUCB using Exp3 [4] to assign probabilities to confidence sets. The analysis of AdaLinUCB combines—in a non-trivial way—the analysis of Exp3 (including a forced exploration term q) with that of SparseLinUCB. Although the resulting regret bound does not improve on OFUL, our algorithm provides a new principled solution to the problem of tuning the radius in (1.1). Experiments show that SparseLinUCB can perform better than OFUL.

1.1 Additional related work

Sparse linear bandits. With the goal of obtaining sparsity-agnostic regret bounds, different types of assumptions on the action set have been considered in the past. Starting from the $\tilde{O}(S\sqrt{T})$ regret upper bound of [18], where the action set is assumed fixed and equal to the hypercube, some works considered stochastic action sets and proved regret bounds depending on spectral parameters of the action distribution, such as the minimum eigenvalue of the covariance matrix [9, 15, 16, 20]. Others assumed a stochastic action set with strong properties, such as compatibility conditions or margin conditions [3, 6, 7, 19, 22]. As far as we know, there has been no research on adaptive adversarial action sets after [2].

Model selection. Sparse linear bandits can be naturally viewed as a bandit model selection problem. For example, Ghosh et al. [14] establish a regret bound of $\tilde{O}(\sqrt{ST} + d^2/\alpha^{4.65})$ for a fixed action set, where α is the minimum absolute value of the nonzero components of θ_* . Quite a bit of work has been devoted to sparse regret bounds in the nested setting. With i.i.d. and fixed-size actions sets, Foster et al. [12] achieve a regret bound of order $\tilde{O}(S^{1/3}T^{2/3}/\gamma^3)$ in the nested setting, where γ is the smallest eigenvalue of the covariance matrix of \mathcal{A}_t . Under the same assumption on the action set, Pacchiano et al. [24, 23] obtain a regret bound of $\tilde{O}(S^2\sqrt{T})$. For adversarial action sets, Cutkosky et al. [8] obtain $\tilde{O}(S^2\sqrt{T})$ in the nested setting. When actions are sampled i.i.d., Cutkosky et al. [8] and Pacchiano et al. [25] obtain a regret bound of $\tilde{O}(S^2\sqrt{T})$ for nested settings. They also obtain simultaneous instance-dependent bounds, in particular, Pacchiano et al. [25] achieve $\tilde{O}((Sd)^2/\Delta)$. Compared to the instance-dependent regret bound, our results are more general, as we allow the action set to be adaptively chosen by an adversary and do not require the nested assumption.

Parameter tuning. Although the theoretical analysis of OFUL only holds for $\gamma_t = O(d \log t)$, smaller choices of the radius in (1.2) are known to perform better in practice. Our design of SparseLinUCB and AdaLinUCB borrows ideas from the parameter tuning setting, which is typically addressed using a set of base algorithms and a randomized master algorithm that adaptively changes the probability of selecting each base algorithm [21, 23, 24]. In particular, AdaLinUCB builds on [11], where they show that running Exp3 as the master algorithm over instances of OFUL with different radii has a better empirical performance than Thompson Sampling and UCB. Yet, they only show a regret bound of $\tilde{O}(T^{2/3})$ when the action set is drawn i.i.d. in each round (they also prove a bound of order \sqrt{T} , but only under additional assumptions on the best model). This is consistent with the results of Pacchiano et al. [24], who also obtained a regret of the same order using Exp3 as master algorithm.

2 Problem definition

In linear bandits, a learner and an adversary interact over T rounds. In each round $t = 1, \dots, T$:

1. The adversary chooses an arm set $\mathcal{A}_t \subset \mathbb{R}^d$;
2. The learner choose an arm $A_t \in \mathcal{A}_t$;
3. The learner obtains a reward X_t .

We assume the adversary is adaptive, i.e., \mathcal{A}_t can depend in an arbitrary way on the (possibly randomized) past choices of the learner. The reward in each round t satisfies

$$X_t = \langle A_t, \theta_* \rangle + \varepsilon_t. \quad (2.1)$$

Here $\theta_* \in \mathbb{R}^d$ is a fixed and unknown target vector and $\{\varepsilon_t\}_{t \in [T]}$ are independent conditionally 1-subgaussian random variables. We also assume $\|\theta_*\|_2 \leq 1$ and $\|a\|_2 \leq 1$ for all $a \in \mathcal{A}_t$ and all $t \in [T]$.¹ The regret of a strategy over T rounds is defined as the difference between the reward obtained by the optimal policy, always choosing the best arm in \mathcal{A}_t , and the reward obtained by the strategy choosing arm $A_t \in \mathcal{A}_t$ for $t \in [T]$,

$$R_T = \sum_{t=1}^T \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle - \sum_{t=1}^T \langle A_t, \theta_* \rangle.$$

¹The choice of the constant 1 is arbitrary. Choosing different constants would scale our bounds similarly to the scaling of the bounds in [2].

In the sparse setting, we would like to devise strategies whose regret depends on

$$S = \|\theta_*\|_0 = \sum_{i=1}^d \mathbb{1}\{\theta_i \neq 0\}$$

corresponding to the number of nonzero components of θ_* .

2.1 Online to confidence set conversions

Establishing a confidence set including the target vector with high probability is at the core of linear bandit algorithms, and our approach for designing a sparsity-agnostic algorithm is based on a result by Abbasi-Yadkori et al. [2]. They show how to construct a confidence set for OFUL [1] based on the predictions of a generic algorithm for online linear regression, a sequential decision-making setting defined as follows. For $t = 1, \dots, T$:

1. The adversary privately chooses input $A_t \in \mathbb{R}^d$ and outcome $X_t \in \mathbb{R}$;
2. The learner observes A_t and chooses prediction $\hat{X}_t \in \mathbb{R}$;
3. The adversary reveals X_t and the learner suffers loss $(\hat{X}_t - X_t)^2$.

The learner's goal in online linear regression is to minimize the following notion of regret against any comparator $\theta \in \mathbb{R}^d$

$$\rho_T(\theta) = \sum_{t=1}^T (X_t - \hat{X}_t)^2 - \sum_{t=1}^T (X_t - \langle A_t, \theta \rangle)^2.$$

The confidence set proposed by Abbasi-Yadkori et al. [2] is established by the following result.

Lemma 2.1 (Abbasi-Yadkori et al. [2, Corollary 2]). *Let $\delta \in (0, 1/4]$ and $\|\theta_*\|_2 \leq 1$. Assume a sequence $\{(A_t, X_t)\}_{t \in [T]}$, where X_t satisfies (2.1) for all $t \in [T]$, is fed to an online linear regression algorithm \mathcal{B} generating predictions $\{\hat{X}_t\}_{t \in [T]}$. Then $\mathbb{P}(\exists t \in [T] : \theta_* \notin \mathcal{C}_t) \leq \delta$, where*

$$\mathcal{C}_t = \left\{ \theta \in \mathbb{R}^d : \|\theta\|_2^2 + \sum_{s=1}^{t-1} (\hat{X}_s - \langle A_s, \theta \rangle)^2 \leq \gamma(\delta, \theta_*) \right\} \quad (2.2)$$

$$\gamma(\delta, \theta_*) = 2 + 2B_T(\theta_*) + 32 \log \left(\frac{\sqrt{8} + \sqrt{1 + B_T(\theta_*)}}{\delta} \right)$$

and $B_T(\theta_*)$ is an upper bound on the regret $\rho_T(\theta_*)$ of \mathcal{B} . When understood from the context, we will abuse the notation and denote the best radius in hindsight by $\gamma(\delta) := \gamma(\delta, \theta_*)$ and the regret bound by B_T .

Gerchinovitz [13] designed an algorithm, SeqSEW, for *sparse* linear regression that bounds $\rho_T(\theta)$ in terms of $\|\theta\|_0$ simultaneously for all comparators $\theta \in \mathbb{R}^d$. Below here, we state his bound in the formulation of Lattimore and Szepesvári [17].

Lemma 2.2 (Lattimore and Szepesvári [17, Theorem 23.6]). *Assume $\max_{t \in [T]} \|A_t\|_2 \leq 1$ and $\max_{t \in [T]} |X_t| \leq 1$. There exists a universal constant c such that algorithm SeqSEW achieves, for any $\theta \in \mathbb{R}^d$,*

$$\rho_T(\theta) \leq B_T(\theta) := c\|\theta\|_0 \left\{ \log(e + T^{1/2}) + C_T \log \left(1 + \frac{\|\theta\|_1}{\|\theta\|_0} \right) \right\}$$

where $C_T = 2 + \log_2 \log(e + T^{1/2})$.

Using the confidence set (2.2) with \mathcal{B} set to SeqSEW, Abbasi-Yadkori et al. [2] achieved the minimax optimal regret bound of $\tilde{O}(\sqrt{SdT})$. However, to construct \mathcal{C}_t , the learner must know $\gamma(\delta)$, which depends on the unknown sparsity level $S = \|\theta_*\|_0$ through B_T .

Algorithm 1 SparseLinUCB

- 1: **Input:** $T \in \mathbb{N}$, $n \in \mathbb{N}$ and $\mathbf{q} \in \Delta_n := \{(q_1, \dots, q_n) \in [0, 1]^n : \sum_{i=1}^n q_i = 1\}$
- 2: **Initialization:** Let $V_0 = I$, $\hat{\theta}_0 = (0, \dots, 0)$
- 3: **for** $t = 1, 2, \dots, T$ **do**
- 4: Receive action set \mathcal{A}_t and draw I_t from distribution \mathbf{q}
- 5: Choose action $A_t = \operatorname{argmax}_{a \in \mathcal{A}_t} \left(\langle a, \hat{\theta}_{t-1} \rangle + \|a\|_{V_{t-1}} \sqrt{2^{I_t} \log T} \right)$
- 6: Receive reward X_t
- 7: $V_t = V_{t-1} + A_t A_t^\top$
- 8: Feed (A_t, X_t) to SeqSEW and obtain prediction \hat{X}_t
- 9: Compute regularized least squares estimate $\hat{\theta}_t = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \left(\|\theta\|_2^2 + \sum_{s=1}^t (\hat{X}_s - \langle \theta, A_s \rangle)^2 \right)$
- 10: **end for**

3 A multi-level sparse linear bandit algorithm

In this section, we introduce our main algorithm, SparseLinUCB, whose pseudo-code is shown in Algorithm 1. The algorithm, which runs SeqSEW as base algorithm \mathcal{B} , uses a hierarchy of confidence sets of increasing radius. In each round $t = 1, \dots, T$, after receiving the action set \mathcal{A}_t , the algorithm draws the index I_t of the confidence set for time t by sampling from the distribution $\mathbf{q} \in \Delta_n := \{(q_1, \dots, q_n) \in [0, 1]^n : \sum_{i=1}^n q_i = 1\}$. Then the algorithm plays the action A_t using the confidence set $\mathcal{C}_t^{I_t} := \{\theta \in \mathbb{R}^d : \|\theta - \theta_{t-1}\|_{V_{t-1}}^2 \leq 2^{I_t} \log T\}$ (where a larger I_t implies a larger radius, and thus more exploration). Following the online to confidence set approach, upon receiving the reward X_t , the algorithm feeds the pair (A_t, X_t) to SeqSEW and uses the prediction \hat{X}_t to update the regularized least squares estimate $\hat{\theta}_t$.

Let $\alpha_i = 2^i \log T$ for all $i \in \mathbb{N}$ and set $n \in \mathbb{N}$ as

$$n = \left\lceil \log_2 \frac{\max_{\theta \in \mathbb{R}^d : \|\theta\|_2 \leq 1} \gamma(1/T, \theta)}{\log T} \right\rceil \quad (3.1)$$

where $\gamma(\delta, \theta)$ is defined in Lemma 2.1 for $\mathcal{B} = \text{SeqSEW}$.

One can check that $n = \Theta(\log d)$ (when $\|\theta\|_0 = d$), which gives $\alpha_n = \Theta(d \log T)$. Our bounds depend on the following quantity, which defines the index of the smallest “safe” confidence set (i.e., the smallest $i \in [n]$ such that $\theta_* \in \mathcal{C}_t^i$ for all $t \in [T]$),

$$o := \operatorname{argmin}_{i \in [n]} \left\{ \gamma(1/T) \leq \alpha_i \right\} \quad (3.2)$$

The choice of our confidence set (Line 4 in Algorithm 1) is justified by the following result, which implies that o is safe.

Lemma 3.1. *For \mathcal{C}_t defined in (2.2), we have that $\mathcal{C}_t \subseteq \mathcal{C}_t^o$ for all $t \in [T]$.*

As we use SeqSEW as base algorithm \mathcal{B} , $\alpha_o = O(S \log T)$. Our main result is an upper bound on the regret of SparseLinUCB.

Theorem 3.2. *The expected regret of SparseLinUCB run with the number of models n in (3.1) and a distribution $\mathbf{q} = \{q_s\}_{s \in [n]}$ satisfies*

$$\mathbb{E}[R_T] = O \left((\log T) \sum_{s \geq o} \sqrt{d 2^s T q_s} + (\log T) \sqrt{S T d / Q} \right)$$

where $Q = \sum_{s \geq o} q_s$.

If the sparsity level S is indeed known, then o in (3.2) can be computed and we get the following bound, which is tight up to log factors [17].

Corollary 3.3. *Assume that the sparsity level S is known and choose the number of models $n > o$ and the distribution $\{q_s\}_{s \in [n]}$ with $q_o = 1$, where o is set as in (3.2). Then, the expected regret of SparseLinUCB is $\mathbb{E}[R_T] = O(\sqrt{SdT} \log T)$.*

An instance-dependent bound. SparseLinUCB also enjoys an instance-dependent regret bound comparable to that of OFUL. Let Δ be the minimum gap between the optimal arm and any suboptimal arms over all rounds,

$$\Delta = \min_{t \in [T]} \min_{a \in \mathcal{A}_t \setminus A_t^*} \langle A_t^* - a, \theta_* \rangle, \quad (3.3)$$

where $A_t^* = \max_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle$ is the optimal arm for round t .

Theorem 3.4. *The expected regret of SparseLinUCB run with the number of models n in (3.1), a distribution $\mathbf{q} = \{q_s\}_{s \in [n]}$ and using SeqSEW as base algorithm satisfies*

$$\mathbb{E}[R_T] = O\left(\frac{(dS/Q) + d^2}{\Delta} (\log T)^2\right)$$

where $Q = \sum_{s \geq o} q_s$.

Sparsity-agnostic tuning of randomization. Next, we look at a specific choice of \mathbf{q} . Fix $C \geq 1$ and let

$$q_s = \begin{cases} C^2 2^{-s} & \text{if } C^2 2^{-s} < 1 \\ \kappa & \text{otherwise,} \end{cases} \quad (3.4)$$

where $\kappa > 0$ is chosen so to normalize the probabilities. It is easy to verify that for any $C \geq 1$,

$$\sum_{s \in [n]} \mathbb{1}\{C^2 2^{-s} < 1\} q_s \leq 1$$

implying that κ can be chosen in $[0, 1]$. Combining Theorem 3.2 and 3.4, we obtain the following corollary providing a hybrid distribution-free and distribution-dependent bound.

Corollary 3.5. *Pick any $C \geq 1$. Let the number of models n as in (3.1) and $\mathbf{q} = \{q_s\}_{s \in [n]}$ be chosen as in (3.4). Then the expected regret of SparseLinUCB is*

$$\mathbb{E}[R_T] = \tilde{O}\left(\min\left\{\max\{C, S/C\}\sqrt{dT}, \frac{\max\{d^2, S^2 d/C^2\}}{\Delta}\right\}\right)$$

For $C = 1$ the above bound is $\tilde{O}(S\sqrt{dT})$, which is tight up to the factor \sqrt{S} due to the lower bound of $\Omega(\sqrt{SdT})$ [17]. However, as mentioned in Lattimore and Szepesvári [17, Section 23.5], no algorithm can enjoy the regret of $\tilde{O}(\sqrt{SdT})$ simultaneously for all possible sparsity levels S . While our worst-case regret bound improves with a smaller S , the problem-dependent regret bound scales at least as $(d^2/\Delta) \log T$, which is independent of S . This raises an interesting question: could the problem-dependent bound also benefit from sparsity? Even with a very small probability p of choosing radius α_n , the expected number of steps using α_n would be pT . The results in [2] demonstrate that running the OFUL algorithm with α_n over pT steps results in a regret of $\tilde{O}(d^2/\Delta)$. One simple way is to decrease the frequency of selecting radius α_n . However, selecting α_n less than d^2/Δ^2 times may prevent the algorithm from obtaining a good enough estimate of θ_* in certain settings.

Remark 3.6. *At first glance, it may seem straightforward to select C in Corollary 3.5, as setting $C = \sqrt{d}$ yields a regret of $\tilde{O}(d^2/\Delta)$ without apparent trade-offs. However, the trade-off lies in balancing the instance-dependent and worst-case regret bounds. Opting for $C = d$ indeed yields an instance-dependent bound of $\tilde{O}(d^2/\Delta)$. However, this comes at the expense of the worst-case bound, which remains $\tilde{O}(d\sqrt{T})$, negating any advantages derived from the sparsity assumption $S \ll d$.*

If the sparsity level S is indeed known, then o in (3.2) can be computed and we get the following bound.

Corollary 3.7. *Assume that the sparsity level S is known and choose $\{q_s\}_{s \in [n]}$ with $q_o = 1$, where o is set as in (3.2). Then, the expected regret of SparseLinUCB is $\mathbb{E}[R_T] = \tilde{O}(\frac{Sd}{\Delta})$.*

We note that by setting $q_o = 1$ in Theorem 3.4, the regret bound becomes $\tilde{O}(d^2/\Delta)$. This result, as detailed in Theorem 3.4, arises from the parameter $q_n > 0$. However, in this case, $q_n = 0$, which allows us to achieve a more favorable regret bound.

Algorithm 2 AdaLinUCB

```

1: Input:  $T \in \mathbb{N}$ ,  $\eta > 0$ ,  $q \in (0, 1]$ 
2: Initialization: Let  $S_{i,0} = 0$  for all  $i \in [n]$ ,  $V_0 = I$ ,  $\hat{\theta}_0 = (0, \dots, 0)$ 
3: for  $t = 1, 2, \dots, T$  do
4:   Receive action set  $\mathcal{A}_t$  and draw a Bernoulli random variable  $Z_t$  with  $\mathbb{P}(Z_t = 1) = q$ 
5:   if  $Z_t = 1$  then
6:     Choose optimistic action  $A_t = \operatorname{argmax}_{a \in \mathcal{A}_t} \left( \langle a, \hat{\theta}_{t-1} \rangle + \|a\|_{V_{t-1}^{-1}} \sqrt{2^n \log T} \right)$ 
7:     Receive reward  $X_t$ ;
8:   else
9:     Draw  $I_t$  from the distribution  $P_{t,i} = \frac{\exp(\eta S_{t-1,i})}{\sum_{j=1}^n \exp(\eta S_{t-1,j})}$  for  $i \in [n]$ ;
10:    Choose action  $A_t = \operatorname{argmax}_{a \in \mathcal{A}_t} \left( \langle a, \hat{\theta}_{t-1} \rangle + \|a\|_{V_{t-1}^{-1}} \sqrt{2^{I_t} \log T} \right)$ 
11:    Receive reward  $X_t$ ;
12:    Compute  $S_{t,j} = S_{t-1,j} - \frac{\mathbb{1}\{I_t = j\}(2 - X_t)/4}{P_{t,j}}$  for  $j \in [n]$ ;
13:   end if
14:    $V_t = V_{t-1} + A_t A_t^\top$ ;
15:   Feed  $(A_t, X_t)$  to SeqSEW and obtain prediction  $\hat{X}_t$ ;
16:   Compute regularized least squares estimate  $\hat{\theta}_t = \operatorname{argmin}_{\theta \in \mathbb{R}^d} \left( \|\theta\|_2^2 + \sum_{s=1}^t (\hat{X}_s - \langle \theta, A_s \rangle)^2 \right)$ ;
17: end for

```

4 Adaptive model selection for stochastic linear bandits

SparseLinUCB is also designed to handle adaptive adversarial action sets. A crucial parameter of SparseLinUCB is $\{q_s\}_{s \in [n]}$, the distribution from which the radius of the confidence set is drawn. It is a natural question whether there exists an algorithm that adaptively updates this distribution based on the observed rewards. In this section we introduce AdaLinUCB (Algorithm 2), which runs Exp3 to dynamically adjust the distribution used by SparseLinUCB.

AdaLinUCB takes as input a forced exploration term q and the learning rate η for Exp3. Similarly to SparseLinUCB, AdaLinUCB designs confidence sets of various radii, but its selection method differs in two aspects. First, with probability q , the algorithm performs exploration based on the confidence set with the largest radius. With probability $1 - q$, the algorithm instead draws the action based on Exp3. The distribution P_t used by Exp3 at round t is based on exponential weights applied to the total estimated loss, denoted by S_t (for technical reasons, we translate losses into rewards). The algorithm then draws I_t from P_t and selects the action A_t based on the confidence set with radius $2^{I_t} \log T$. Finally, reward X_t is observed and the pair (A_t, X_t) is fed to SeqSEW. The prediction \hat{X}_t returned by SeqSEW is used to update the regularized least squares estimate $\hat{\theta}_t$.

The following theorem states the theoretical regret upper bound of AdaLinUCB.

Theorem 4.1. *If the random independent noise ε_t in (2.1) satisfies $\varepsilon_t \in [-1, 1]$ for all $t \in [T]$, then the regret of AdaLinUCB run with $\eta = \sqrt{(\log n)/(Tn)}$ for n in (3.1) and $q \in (0, 1]$ satisfies*

$$\begin{aligned} \mathbb{E}[R_T] &\leq \left(\sqrt{8\alpha_n q} + 4\sqrt{(2\alpha_o)/q} \right) \sqrt{dT \log \left(1 + \frac{TL^2}{d} \right) + 1} + O(\sqrt{nT \log n}) \\ &= \tilde{O} \left(\max \left\{ \sqrt{dq}, \sqrt{S/q} \right\} \sqrt{dT} \right). \end{aligned}$$

Although AdaLinUCB dynamically adjusts the distribution used by SparseLinUCB and may achieve better empirical performance, its regret bound is no better than that of SparseLinUCB. The issue is that the action chosen by AdaLinUCB in Line 10 does not ensure enough exploration to control the regret. Consequently, the algorithm needs to choose the optimistic action in Line 6 with constant probability q . SparseLinUCB has a similar parameter, Q , that bounds from the above the probability

of choosing the optimistic action. The key difference is that Q can be optimized for an unknown S by carefully selecting the distribution $\mathbf{q} = \{q_s\}_{s \geq 1}$, whereas the parameter q does not provide a similar flexibility.

5 Model selection experiments

In this section we describe some experiments we performed on synthetic data to verify whether AdaLinUCB could outperform OFUL in a model selection task. We also test the empirical performance of SparseLinUCB on the same data (additional details on all the algorithms and the experimental setting are in Appendix E). The data for our model selection experiments are generated using targets θ_* with different sparsity levels, as we know that sparsity affects the radius of the optimal confidence set. On the other hand, since no efficient implementation of SeqSEW is known [17, Section 23.5], we cannot implement the online to confidence set approach as described in [2] to capture sparsity. Instead, we run SparseLinUCB and AdaLinUCB with $\hat{X}_t = X_t$ for all $t \in [T]$, which—due to the form of our confidence sets—amounts to running the algorithms over multiple instances of OFUL with different choices of radius α_i for $i \in [n]$.

We run SparseLinUCB and AdaLinUCB with $\alpha_i = \alpha_{i,t} = 2^i \log t$ (a mildly time-dependent choice) for $i = 0, 1, \dots, \log_2 d$. We also include $\alpha_0 = 0$ corresponding to the greedy strategy $A_t = \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a, \hat{\theta}_{t-1} \rangle$. The suffix `_Unif` indicates $\{q_s\}_{s \in [n]}$ set to $(\frac{1}{n}, \dots, \frac{1}{n})$. The suffix `_Theory` indicates $q_s = \Theta(2^{-s})$ for $s = 0, \dots, n$ as prescribed by (3.4). Finally, we included SparseLinUCB_Known using $q_i = \mathbb{1}\{i = o\}$ to test the performance when the optimal index o (for the given S) is known in advance (see Corollary 3.3). We run our experiments with a fixed set of random actions, $\mathcal{A}_t = \mathcal{A}$ for all $t \in [T]$, where $|\mathcal{A}| = 30$ and \mathcal{A} is a set of vectors drawn i.i.d. from the unit sphere in \mathbb{R}^{16} . The target vector θ_* is a S -sparse ($S = 1, 2, 4, 8, 16$) vector whose non-zero coordinates are drawn from the unit sphere in \mathbb{R}^S . The noise ε_t is drawn i.i.d. from the uniform distribution over $[-1, 1]$. Each curve is an average over 20 repetitions with $T = 10^4$ where, in each repetition, we draw fresh instances of \mathcal{A} and θ_* .

As our implementations are not sparsity-aware, we cannot expect the regret to strongly depend on the sparsity level. Indeed, only the regret of SparseLinUCB_Known (which is tuned to the sparsity S) is significantly affected by sparsity. The theory-driven choice of $\{q_s\}_{s \in [n]}$ (SparseLinUCB_Theory) performs better than the uniform assignment (SparseLinUCB_Unif), and is in the same ballpark as OFUL. On the other hand, AdaLinUCB_Unif and AdaLinUCB_Theory outperform all the competitors, including OFUL. This provides evidence that using Exp3 for adaptive model selection may significantly boost the empirical performance of stochastic linear bandits.

6 Limitations and open problems

Unlike previous works, we prove sparsity-agnostic regret bounds with no assumptions on the action sets or on θ_* (other than boundedness of $\|\theta_*\|$ and $\|a\|$ for $a \in \mathcal{A}_t$, which are rather standard assumptions). For AdaLinUCB, however, we do require boundedness of the noise ε_t (instead of just subgaussianity). We conjecture this requirement could be dropped at the expense of a further $\log T$ factor in the regret. Finally, for efficiency reasons our implementations are not designed to capture sparsity. Hence our experiments are limited to testing the impact of model selection.

Our work leaves some open problems:

1. Proving a lower bound on the regret of sparse linear bandits when the sparsity level is unknown to the learner would be important. Citing again [17, Section 23.5], no algorithm can enjoy regret $\tilde{O}(\sqrt{SdT})$ simultaneously for all sparsity levels S . However, we do not know whether the known lower bound $\Omega(\sqrt{SdT})$ can be strengthened to $\Omega(S\sqrt{dT})$ in the agnostic case.
2. Our instance-dependent regret bound is of order $\tilde{O}(\max\{S^2, d\} \frac{d}{\Delta})$. It would be interesting to prove an upper bound that improves on the factor d^2/Δ , or a lower bound showing that d^2/Δ cannot be improved on.
3. Our bound on the regret of AdaLinUCB looks pessimistic due to the presence of the constant exploration probability q . It would be interesting to prove a bound that more closely reflects the good empirical performance of this algorithm.

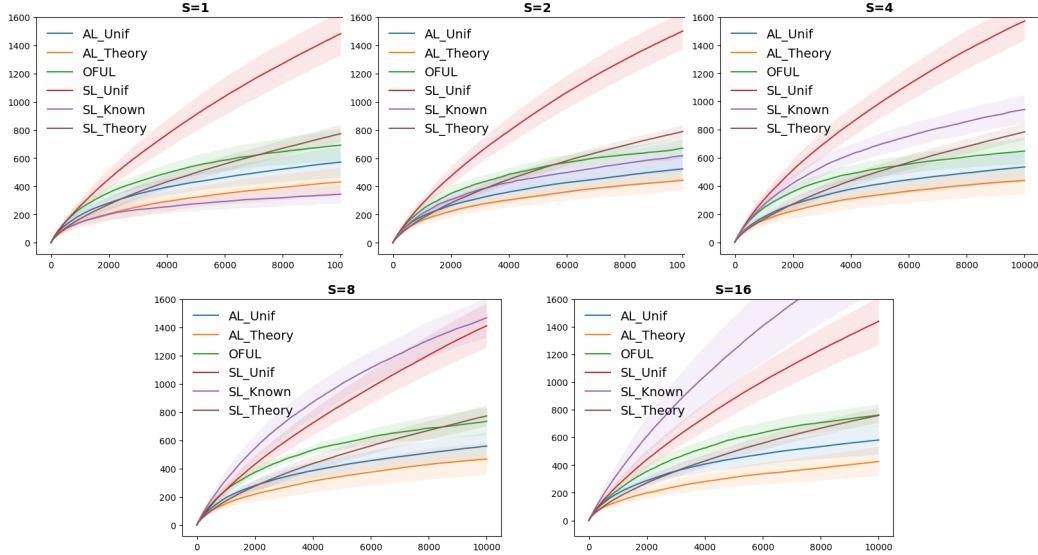


Figure 1: Experimental results with different sparsity levels $S \in \{1, 2, 4, 8, 16\}$. In each plot, the X-axis are time steps in $[1, 10^4]$ and the Y-axis is cumulative regret. AL stands for AdaLinUCB and SL stands for SparseLinUCB.

Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers for their helpful comments. This research is supported by the MUR PRIN grant 2022EKNE5K (Learning in Markets and Society), the FAIR (Future Artificial Intelligence Research) project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme, the EU Horizon CL4-2022-HUMAN-02 research, innovation action under grant agreement 101120237, project ELIAS (European Lighthouse of AI for Sustainability), a Singapore Ministry of Education AcRF Tier 2 grant (A-8000423-00-00), and the National Research Foundation, Singapore under its AI Singapore Program (AISG Award No: AISG-PhD/2021-01-004[T]).

In particular, NCB and KJ acknowledge the financial support from the MUR PRIN grant, the FAIR project, and the ELIAS project. TJ is funded by a Singapore Ministry of Education AcRF Tier 2 grant (A-8000423-00-00) and the National Research Foundation, Singapore under its AI Singapore Program (AISG Award No: AISG-PhD/202101-004[T]).

References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- [2] Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9. PMLR, 2012.
- [3] Kaito Ariu, Kenshi Abe, and Alexandre Proutière. Thresholded lasso bandit. In *International Conference on Machine Learning*, pages 878–928. PMLR, 2022.
- [4] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- [5] Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- [6] Sunrit Chakraborty, Saptarshi Roy, and Ambuj Tewari. Thompson sampling for high-dimensional sparse linear contextual bandits. In *International Conference on Machine Learning*, pages 3979–4008. PMLR, 2023.

[7] Yi Chen, Yining Wang, Ethan X Fang, Zhaoran Wang, and Runze Li. Nearly dimension-independent sparse linear bandit over small action spaces via best subset selection. *Journal of the American Statistical Association*, 119(545):246–258, 2024.

[8] Ashok Cutkosky, Christoph Dann, Abhimanyu Das, Claudio Gentile, Aldo Pacchiano, and Manish Purohit. Dynamic balancing for model selection in bandits and rl. In *International Conference on Machine Learning*, pages 2276–2285. PMLR, 2021.

[9] Yan Dai, Ruosong Wang, and Simon Shaolei Du. Variance-aware sparse linear bandits. In *The Eleventh International Conference on Learning Representations*, 2023.

[10] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. 2008.

[11] Qin Ding, Yue Kang, Yi-Wei Liu, Thomas Chun Man Lee, Cho-Jui Hsieh, and James Sharpnack. Syndicated bandits: A framework for auto tuning hyper-parameters in contextual bandit algorithms. *Advances in Neural Information Processing Systems*, 35:1170–1181, 2022.

[12] Dylan J Foster, Akshay Krishnamurthy, and Haipeng Luo. Model selection for contextual bandits. *Advances in Neural Information Processing Systems*, 32, 2019.

[13] Sébastien Gerchinovitz. Sparsity regret bounds for individual sequences in online linear regression. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 377–396. JMLR Workshop and Conference Proceedings, 2011.

[14] Avishek Ghosh, Abishek Sankararaman, and Ramchandran Kannan. Problem-complexity adaptive model selection for stochastic linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 1396–1404. PMLR, 2021.

[15] Botao Hao, Tor Lattimore, and Mengdi Wang. High-dimensional sparse linear bandits. *Advances in Neural Information Processing Systems*, 33:10753–10763, 2020.

[16] Kyoungseok Jang, Chicheng Zhang, and Kwang-Sung Jun. Popart: Efficient sparse regression and experimental design for optimal sparse linear bandits. *Advances in Neural Information Processing Systems*, 35:2102–2114, 2022.

[17] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

[18] Tor Lattimore, Koby Crammer, and Csaba Szepesvári. Linear multi-resource allocation with semi-bandit feedback. *Advances in Neural Information Processing Systems*, 28, 2015.

[19] Ke Li, Yun Yang, and Naveen N Narisetty. Regret lower bound and optimal algorithm for high-dimensional contextual linear bandit. *Electronic Journal of Statistics*, 15(2):5652–5695, 2021.

[20] Wenjie Li, Adarsh Barik, and Jean Honorio. A simple unified framework for high dimensional bandit problems. In *International Conference on Machine Learning*, pages 12619–12655. PMLR, 2022.

[21] Teodor Vanislavov Marinov and Julian Zimmert. The pareto frontier of model selection for general contextual bandits. *Advances in Neural Information Processing Systems*, 34:17956–17967, 2021.

[22] Min-Hwan Oh, Garud Iyengar, and Assaf Zeevi. Sparsity-agnostic lasso bandit. In *International Conference on Machine Learning*, pages 8271–8280. PMLR, 2021.

[23] Aldo Pacchiano, Christoph Dann, Claudio Gentile, and Peter Bartlett. Regret bound balancing and elimination for model selection in bandits and RL. *arXiv preprint arXiv:2012.13045*, 2020.

[24] Aldo Pacchiano, My Phan, Yasin Abbasi Yadkori, Anup Rao, Julian Zimmert, Tor Lattimore, and Csaba Szepesvari. Model selection in contextual stochastic bandit problems. *Advances in Neural Information Processing Systems*, 33:10328–10337, 2020.

[25] Aldo Pacchiano, Christoph Dann, and Claudio Gentile. Best of both worlds model selection. *Advances in Neural Information Processing Systems*, 35:1883–1895, 2022.

[26] Vidyashankar Sivakumar, Steven Wu, and Arindam Banerjee. Structured linear contextual bandits: A sharp and geometric smoothed analysis. In *International Conference on Machine Learning*, pages 9026–9035. PMLR, 2020.

Table 2: Notation.

Symbol	Description
\mathcal{A}_t	Action set at time t
θ_*	True parameter of the linear model
d	Dimension of θ_*
S	Number of nonzero components in θ_*
X_t	Random reward of pulling arm $A_t \in \mathcal{A}_t$
\hat{X}_t	Prediction of SeqSEW at time t
\mathcal{C}_t^s	The confidence set with radius $2^s \log T$ and center $\hat{\theta}_{t-1}$
I_t	The index of the confidence set drawn at time t
q_s	The probability $\mathbb{P}(I_t = s)$ of drawing index s of the confidence set in SparseLinUCB
o	The index of the smallest safe confidence set, defined in (3.2)
Q	The probability $\mathbb{P}(I_t \geq o) = q_o + \dots + q_n$ of drawing a safe confidence set in SparseLinUCB
A_t^s	The optimistic action for \mathcal{C}_t^s
A_t^*	The optimal action in \mathcal{A}_t
\mathcal{E}	Event that $\theta_* \in \mathcal{C}_t$ for all $t \in [T]$
$\det V$	Determinant of V

A Notation

In Table 2 we list the most used notations. Next, we recall some definitions that are used throughout this appendix. We have $V_t = I + \sum_{s=1}^t A_s A_s^\top$ and

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

where $A_s \in \mathcal{A}_s$ is the action chosen by the learner at round t . Recall that $\alpha_i = 2^i \log T$. For $i \in [n]$,

$$A_t^i = \operatorname{argmax}_{a \in \mathcal{A}_t} \max_{\theta \in \mathcal{C}_t^i} \langle \theta, a \rangle = \operatorname{argmax}_{a \in \mathcal{A}_t} \left(\langle a, \hat{\theta}_{t-1} \rangle + \|a\|_{V_{t-1}^{-1}} \sqrt{\alpha_i} \right)$$

where

$$\mathcal{C}_t^i = \left\{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_{t-1}\|_{V_{t-1}}^2 \leq \alpha_i \right\}.$$

Finally, recall definition (2.2) with $\delta = 1/T$,

$$\mathcal{C}_t = \left\{ \theta \in \mathbb{R}^d : \|\theta\|_2^2 + \sum_{s=1}^{t-1} (\hat{X}_s - \langle A_s, \theta \rangle)^2 \leq \gamma(1/T) \right\}.$$

and recall that $\mathcal{E} = \bigcap_{t \in [T]} \{\theta_* \in \mathcal{C}_t\}$.

B Analysis of SparseLinUCB

Theorem 3.2. *The expected regret of SparseLinUCB run with the number of models n in (3.1) and a distribution $\mathbf{q} = \{q_s\}_{s \in [n]}$ satisfies*

$$\mathbb{E}[R_T] = O \left((\log T) \sum_{s \geq o} \sqrt{d 2^s T q_s} + (\log T) \sqrt{S T d / Q} \right)$$

where $Q = \sum_{s \geq o} q_s$.

Proof. Lemma 3.1 implies $\mathcal{C}_t \subset \mathcal{C}_t^o$. Hence, if \mathcal{E} is true and $s < o$, then

$$\begin{aligned} \langle \theta_*, A_t^s \rangle &\geq \langle \hat{\theta}_{t-1}, A_t^s \rangle - \sqrt{\alpha_o} \|A_t^s\|_{V_{t-1}^{-1}} & (\theta_* \in \mathcal{C}_t \subset \mathcal{C}_t^o) \\ &= \left(\langle \hat{\theta}_{t-1}, A_t^s \rangle + \sqrt{\alpha_s} \|A_t^s\|_{V_{t-1}^{-1}} \right) - (\sqrt{\alpha_o} + \sqrt{\alpha_s}) \|A_t^s\|_{V_{t-1}^{-1}} \end{aligned}$$

$$\begin{aligned}
&\geq \left(\langle \hat{\theta}_{t-1}, A_t^o \rangle + \sqrt{\alpha_s} \|A_t^o\|_{V_{t-1}^{-1}} \right) - (\sqrt{\alpha_o} + \sqrt{\alpha_s}) \|A_t^s\|_{V_{t-1}^{-1}} \quad (\text{maximality of } A_t^s \text{ in } \mathcal{C}_t^s) \\
&= \langle \hat{\theta}_{t-1}, A_t^o \rangle + \sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} - (\sqrt{\alpha_o} - \sqrt{\alpha_s}) \|A_t^o\|_{V_{t-1}^{-1}} - (\sqrt{\alpha_o} + \sqrt{\alpha_s}) \|A_t^s\|_{V_{t-1}^{-1}} \\
&\geq \langle \theta_*, A_t^o \rangle - (\sqrt{\alpha_o} - \sqrt{\alpha_s}) \|A_t^o\|_{V_{t-1}^{-1}} - (\sqrt{\alpha_o} + \sqrt{\alpha_s}) \|A_t^s\|_{V_{t-1}^{-1}} \quad (\theta_* \in \mathcal{C}_t \subset \mathcal{C}_t^o) \\
&\geq \langle \theta_*, A_t^o \rangle - (\sqrt{\alpha_o} - \sqrt{\alpha_s}) \|A_t^o\|_{V_{t-1}^{-1}} - (\sqrt{\alpha_o} + \sqrt{\alpha_s}) \|A_t^o\|_{V_{t-1}^{-1}} \quad (\text{by Lemma D.3}) \\
&= \langle \theta_*, A_t^o \rangle - 2\sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} \\
&\geq \langle \theta_*, A_t^* \rangle - 3\sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}}
\end{aligned} \tag{B.1}$$

where in the first and the third inequalities, we use the fact that for any $A \in \mathbb{R}^d$,

$$\begin{aligned}
\langle \theta_* - \hat{\theta}_{t-1}, A \rangle &\leq \|\theta_* - \hat{\theta}_{t-1}\|_{V_{t-1}} \|A\|_{V_{t-1}^{-1}}, & (\text{Cauchy-Schwarz inequality}) \\
&\leq \sqrt{\alpha_o} \|A\|_{V_{t-1}^{-1}} & (\text{due to } \theta_* \in \mathcal{C}_t^o)
\end{aligned}$$

and for the last inequality,

$$\langle \theta_*, A_t^* \rangle \leq \langle \hat{\theta}_{t-1}, A_t^o \rangle + \sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} \quad (\theta_* \in \mathcal{C}_t^o, \text{ maximality of } A_t^o).$$

We can decompose the regret as follows,

$$\begin{aligned}
R_T &= \sum_{t=1}^T \mathbb{1}\{\mathcal{E}\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle \\
&= \sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \sum_{t=1}^T \mathbb{1}\{I_t < o, \mathcal{E}\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle \\
&\leq \sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle \\
&\quad + \sum_{t=1}^T \mathbb{1}\{I_t < o, \mathcal{E}\} \min \left\{ 2, 3\sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} \right\} \\
&\leq \sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \underbrace{\sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle}_{\spadesuit} + \underbrace{\sum_{t=1}^T \left\{ 2, 3\sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} \right\}}_{\clubsuit}.
\end{aligned} \tag{B.2}$$

Since $\mathbb{P}(\mathcal{E}^c) \leq \delta \leq \frac{1}{T}$, the first sum in the above line is easily bounded,

$$\sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle \leq 2T\mathbb{P}(\mathcal{E}^c) \leq 2.$$

Bounding term \clubsuit . For each $s \geq 0$, let

$$\mathcal{T}^s = \left\{ t \in [T] : \det(V_{t-1}) \in \left[2^{sd}, 2^{(s+1)d} \right) \right\}.$$

Note that $\det(V_t)$ is monotone increasing w.r.t t . Define $s' = \lceil \log_2 \det(V_T)/d \rceil$. Then,

$$[T] = \bigcup_{s=1}^{s'} \mathcal{T}^s.$$

Therefore,

$$\clubsuit \leq \sum_{s=1}^{s'} \sum_{t \in \mathcal{T}^s} \min \left\{ 2, 3\sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} \right\}.$$

By applying Lemma D.5, we obtain

$$\begin{aligned}
\mathbb{E}[\clubsuit] &= \mathbb{E} \left[\sum_{s=1}^{s'} \sum_{t \in \mathcal{T}^s} \min \left\{ 2, 3\sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} \right\} \right] \\
&= 3\sqrt{\alpha_o} \mathbb{E} \left[\sum_{s=1}^{s'} \sum_{t \in \mathcal{T}^s} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}} \right\} \right] \quad (\alpha_o \geq 1) \\
&\leq 3\sqrt{\alpha_o} \sqrt{T \cdot \mathbb{E} \left[\sum_{s=1}^{s'} \left[\sum_{t \in \mathcal{T}^s} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} \right] \right]} \quad (\text{Cauchy-Schwarz inequality}) \\
&\leq 3\sqrt{\alpha_o T} \sqrt{\sum_{s=1}^{s'} (2d/Q + 1/Q)} \quad (\text{due to Lemma D.5}) \\
&\leq 3\sqrt{\alpha_o T} \sqrt{(2d+1)s'/Q} \\
&= 3\sqrt{\alpha_o T} \sqrt{(2d+1)/Q \lceil \log_2 \det(V_T)/d \rceil} \\
&= O\left(\log T \sqrt{STd/Q}\right).
\end{aligned}$$

Bounding term ♠. Let $T_s = \sum_{t=1}^T \mathbb{1}\{I_t = s\}$.

$$\begin{aligned}
\mathbb{E}[\spadesuit] &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle \right] \\
&\leq \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \cdot \min \left\{ 2, \sqrt{\alpha_{I_t}} \|A_t^{I_t}\|_{V_{t-1}^{-1}} \right\} \right] \\
&\leq 2 \sum_{s \geq o} \mathbb{E} \left[\sqrt{\alpha_s T_s} \sqrt{\sum_{t \in [T]} \mathbb{1}\{I_t = s\} \min \left\{ 1, \|A_t^s\|_{V_{t-1}^{-1}}^2 \right\}} \right] \quad (\text{Cauchy-Schwarz inequality}) \\
&\leq 2 \sum_{s \geq o} \mathbb{E} \left[\sqrt{\alpha_s T_s} \sqrt{\sum_{t \in [T]} \min \left\{ 1, \|A_t\|_{V_{t-1}^{-1}}^2 \right\}} \right] \\
&\leq 2 \sum_{s \geq o} \mathbb{E} \left[\sqrt{\alpha_s T_s} \sqrt{2 \log \det V_T} \right] \quad (\text{Lemma D.1}) \\
&\leq 2 \sum_{s \geq o} \mathbb{E} \left[\sqrt{\alpha_s T_s} \cdot O(\sqrt{d \log T}) \right] \quad (\text{upper bound on } \det(V_T)) \\
&\leq 2 \sum_{s \geq o} \sqrt{\alpha_s} \sqrt{\mathbb{E}[T_s]} \cdot O(\sqrt{d \log T}) \quad (\text{Jensen's inequality}) \\
&= O\left(\sum_{s \geq o} \sqrt{\alpha_s d T q_s \log T}\right) \quad (\text{because } \mathbb{P}(I_t = s) = q_s).
\end{aligned}$$

Substituting the bounds of ♠ and ♣ to (B.2), we have

$$\mathbb{E}[R_T] \leq O\left(\sum_{s \geq o} \sqrt{d \alpha_s T \log T q_s} + \log T \sqrt{SdT/Q}\right)$$

concluding the proof. \square

Theorem 3.4. *The expected regret of SparseLinUCB run with the number of models n in (3.1), a distribution $\mathbf{q} = \{q_s\}_{s \in [n]}$ and using SeqSEW as base algorithm satisfies*

$$\mathbb{E}[R_T] = O\left(\frac{(dS/Q) + d^2}{\Delta} (\log T)^2\right)$$

where $Q = \sum_{s \geq o} q_s$.

Proof.

$$\begin{aligned}
R_T &= \sum_{t=1}^T \mathbb{1}\{\mathcal{E}\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle \\
&\leq \sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \frac{\langle \theta_*, A_t^* - A_t^{I_t} \rangle^2}{\Delta} \\
&\quad + \sum_{t=1}^T \mathbb{1}\{I_t < o, \mathcal{E}\} \frac{\langle \theta_*, A_t^* - A_t^{I_t} \rangle^2}{\Delta} \quad (\text{minimality of } \Delta) \\
&\leq \sum_{t=1}^T \mathbb{1}\{\mathcal{E}^c\} \langle \theta_*, A_t^* - A_t^{I_t} \rangle + \underbrace{\sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \frac{\langle \theta_*, A_t^* - A_t^{I_t} \rangle^2}{\Delta}}_{\spadesuit} \\
&\quad + \underbrace{9 \sum_{t=1}^T \mathbb{1}\{I_t < o, \mathcal{E}\} \frac{\min\{1, \alpha_o \|A_t^o\|_{V_{t-1}}^2\}}{\Delta}}_{\clubsuit} \quad (\text{by applying (B.1) to } \clubsuit)
\end{aligned}$$

Bounding term ♣. Let $s' = \lceil \log_2 \det(V_T)/d \rceil$. By applying Lemma D.5, we obtain

$$\begin{aligned}
\mathbb{E}[\clubsuit] &\leq \frac{9}{\Delta} \mathbb{E} \left[\sum_{s \geq 0} \sum_{t \in \mathcal{T}^s} \min \left\{ 1, \alpha_o \|A_t^o\|_{V_{t-1}}^2 \right\} \right] \\
&\leq \frac{9\alpha_o}{\Delta} \sum_{s=1}^{s'} \mathbb{E} \left[\sum_{t \in \mathcal{T}^s} \min \left\{ 1, \|A_t^o\|_{V_{t-1}}^2 \right\} \right] \quad (\alpha_o \geq 1) \\
&\leq \frac{9\alpha_o}{\Delta} \sum_{s=1}^{s'} (2d/Q + 1/Q) \quad (\text{due to Lemma D.5}) \\
&\leq \frac{27\alpha_o ds'/Q}{\Delta} \\
&= O \left(\frac{dS(\log T)^2/Q}{\Delta} \right) \quad (\text{because } s' = O(\log T).)
\end{aligned}$$

Bounding term ♠.

$$\begin{aligned}
\mathbb{E}[\spadesuit] &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \frac{\langle \theta_*, A_t^* - A_t^{I_t} \rangle^2}{\Delta} \right] \\
&\leq \frac{4}{\Delta} \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \cdot \min \left\{ 1, \alpha_{I_t} \|A_t^{I_t}\|_{V_{t-1}}^2 \right\} \right] \\
&\leq 4 \sum_{s \geq o} \frac{\alpha_s}{\Delta} \mathbb{E} \left[\sum_{t \in [T]} \min \left\{ 1, \|A_t\|_{V_{t-1}}^2 \right\} \right] \\
&\leq 8 \sum_{s \geq o} \frac{\alpha_s}{\Delta} \cdot \log \det V_T \quad (\text{due to Lemma D.1}) \\
&= O \left(\frac{d^2(\log T)^2}{\Delta} \right) \quad (\text{B.3})
\end{aligned}$$

where the first inequality comes from

$$\langle \theta_*, A_t^* - A_t^{I_t} \rangle \leq \max_{\theta \in \mathcal{C}_{t-1}^{I_t}} \langle \theta - \theta_*, A_t^{I_t} \rangle$$

$$\begin{aligned}
&\leq \max_{\theta \in \mathcal{C}_{t-1}^{I_t}} \|\theta - \theta_*\|_{V_{t-1}} \|A_t^{I_t}\|_{V_{t-1}^{-1}} \quad && \text{(Cauchy-Schwarz Inequality)} \\
&\leq 2\sqrt{\alpha_{I_t}} \|A_t^{I_t}\|_{V_{t-1}^{-1}}
\end{aligned}$$

where the last inequality holds because \mathcal{E} and $I_t \geq o$ both hold, and so $\theta_* \in \mathcal{C}_{t-1} \subset \mathcal{C}_{t-1}^{I_t}$ due to Lemma 3.1. The factor 2 is due to an application of the triangular inequality. Substituting the bounds of \spadesuit and \clubsuit in (B.2), we have

$$\mathbb{E}[R_T] = O\left(\max\{S/Q, d\} \cdot \frac{d(\log T)^2}{\Delta}\right).$$

□

Corollary B.1. *Pick any $C \geq 1$ and let $\{q_s\}_{s \in [n]}$ be chosen as in (3.4). Then the expected regret of SparseLinUCB is*

$$\mathbb{E}[R_T] = \tilde{O}\left(\min\left\{\max\{C, S/C\}\sqrt{dT}, \frac{\max\{d^2, S^2d/C^2\}}{\Delta}\right\}\right)$$

Proof. We consider the following cases based on the value of C .

Case 1: $C^2 < 2^o$.

$q_o = C^2 \cdot 2^{-o}$. According to Theorem 3.2,

$$\begin{aligned}
\mathbb{E}[R_T] &= O\left((\log T) \sum_{s \geq o} \sqrt{d2^s T q_s} + (\log T) \sqrt{STd/Q}\right) \\
&\leq O\left((\log T) n \sqrt{d2^o T q_o} + (\log T) \sqrt{STd/q_o}\right) \\
&= \tilde{O}\left(\max\{C, S/C\}\sqrt{dT}\right). \quad \text{(due to } n = O(\log d) \text{ and } 2^o = \Theta(S \log T)\text{)}
\end{aligned}$$

Besides, according to Theorem 3.4,

$$\begin{aligned}
\mathbb{E}[R_T] &= O\left(\max\{S/Q, d\} \cdot \frac{d \log^2 T}{\Delta}\right) \\
&= O\left(\max\{S/q_o, d\} \cdot \frac{d \log^2 T}{\Delta}\right) \\
&= \tilde{O}\left(\frac{\max\{d^2, S^2d/C^2\}}{\Delta}\right).
\end{aligned}$$

Therefore,

$$\mathbb{E}[R_T] = \tilde{O}\left(\min\left\{\max\{C, S/C\}\sqrt{dT}, \frac{\max\{d^2, S^2d/C^2\}}{\Delta}\right\}\right).$$

Case 2: $C^2 \in [2^o, 2^n]$.

For s with $C^2 < 2^s$, $q_s = C^2 2^{-s}$. Let $o' = \arg \min_{s > o} \{C^2 < 2^s\}$. Then, $q_{o'} \geq 1/4$. According to Theorem 3.2,

$$\begin{aligned}
\mathbb{E}[R_T] &= O\left((\log T) \sum_{s \geq o} \sqrt{d2^s T q_s} + (\log T) \sqrt{STd/Q}\right) \\
&\leq O\left((\log T) \sum_{s \in [o, o')} \sqrt{d2^s T} + (\log T) \sum_{s \in [o', n]} \sqrt{d2^s T q_s} + (\log T) \sqrt{STd/q_{o'}}\right) \\
&\leq \tilde{O}\left(\sqrt{d2^{o'} T} + nC\sqrt{dT} + \sqrt{STd}\right)
\end{aligned}$$

$$= \tilde{O} \left(\max\{C, S/C\} \sqrt{dT} \right). \quad (\text{due to } C^2 = \Theta(2^{o'}) \text{ and } 2^{o'} \geq 2^o \geq S)$$

Besides, according to Theorem 3.4,

$$\begin{aligned} \mathbb{E}[R_T] &= O \left(\max\{S/Q, d\} \cdot \frac{d(\log T)^2}{\Delta} \right) \\ &= O \left(\max\{S/q_{o'}, d\} \cdot \frac{d(\log T)^2}{\Delta} \right) \\ &= \tilde{O} \left(\frac{\max\{d^2, S^2 d/C^2\}}{\Delta} \right). \end{aligned}$$

Therefore,

$$\mathbb{E}[R_T] = \tilde{O} \left(\min \left\{ \max \{C, S/C\} \sqrt{dT}, \frac{\max \{d^2, S^2 d/C^2\}}{\Delta} \right\} \right).$$

Case $C^2 \geq 2^n$: $q_s = 1/n$ for all $s \in [n]$. It is easy to verify that $\mathbb{E}[R_T] = \tilde{O}(d\sqrt{T})$ and $\mathbb{E}[R_T] = \tilde{O}(d^2/\Delta)$. Thus,

$$\mathbb{E}[R_T] = \tilde{O} \left(\min \left\{ \max \{C, S/C\} \sqrt{dT}, \frac{\max \{d^2, S^2 d/C^2\}}{\Delta} \right\} \right).$$

□

Corollary B.2. Assume that the sparsity level S is known and choose a distribution $\mathbf{q} = \{q_s\}_{s \in [n]}$ with $q_o = 1$, where o is set as in (3.2). Then, the expected regret of SparseLinUCB is $\mathbb{E}[R_T] = \tilde{O}(\frac{Sd}{\Delta})$.

Proof. We follow the same steps as in the proof of Theorem 3.4. The only difference lies in the bounding term ♠ in (B.3). We have

$$\begin{aligned} \mathbb{E}[\spadesuit] &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}\{I_t \geq o, \mathcal{E}\} \frac{\langle \theta_*, A_t^* - A_t^{I_t} \rangle^2}{\Delta} \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}\{\mathcal{E}\} \frac{\langle \theta_*, A_t^* - A_t \rangle^2}{\Delta} \right] \\ &\leq \frac{4\alpha_o}{\Delta} \mathbb{E} \left[\sum_{t \in [T]} \min \left\{ 1, \|A_t\|_{V_{t-1}}^2 \right\} \right] \\ &= O \left(\frac{Sd(\log T)^2}{\Delta} \right) \end{aligned} \tag{B.4}$$

where the second equality is because $q_o = 1$ and so $A_t^{I_t} = A_t^o = A_t$, and the first inequality comes from

$$\begin{aligned} \langle \theta_*, A_t^* - A_t \rangle &\leq \max_{\theta \in \mathcal{C}_{t-1}^o} \langle \theta - \theta_*, A_t \rangle \\ &\leq \max_{\theta \in \mathcal{C}_{t-1}^o} \|\theta - \theta_*\|_{V_{t-1}} \|A_t^o\|_{V_{t-1}} \quad (\text{Cauchy-Schwarz Inequality}) \\ &\leq 2\sqrt{\alpha_o} \|A_t\|_{V_{t-1}} \end{aligned}$$

where the last inequality holds because \mathcal{E} and $I_t = o$ both hold, and so $\theta_* \in \mathcal{C}_{t-1} \subset \mathcal{C}_{t-1}^o$ due to Lemma 3.1. Substituting the bounds of ♠ and ♣ in (B.2), we have

$$\mathbb{E}[R_T] = O \left(\frac{Sd(\log T)^2}{\Delta} \right).$$

□

C Analysis of AdaLinUCB

Theorem 4.1. *If the random independent noise ε_t in (2.1) satisfies $\varepsilon_t \in [-1, 1]$ for all $t \in [T]$, then the regret of AdaLinUCB run with $\eta = \sqrt{(\log n)/(Tn)}$ for n in (3.1) and $q \in (0, 1]$ satisfies*

$$\begin{aligned}\mathbb{E}[R_T] &\leq \left(\sqrt{8\alpha_n q} + 4\sqrt{(2\alpha_o)/q} \right) \sqrt{dT \log \left(1 + \frac{TL^2}{d} \right)} + 1 + O(\sqrt{nT \log n}) \\ &= \tilde{O} \left(\max \left\{ \sqrt{dq}, \sqrt{S/q} \right\} \sqrt{dT} \right).\end{aligned}$$

Proof. Let \mathcal{T}_1 be the set of time steps where $Z_t = 1$ in Line 4 of AdaLinUCB and let $\mathcal{T}_2 = [T] \setminus \mathcal{T}_1$. For all $t \in \mathcal{T}_2$ the action $A_t \in \mathcal{A}_t$ is chosen by Exp3 through an adversarial mapping $\mu_t : [n] \rightarrow \mathcal{A}_t$ defined in Lines 9–10. The resulting reward $X_t \in [-2, 2]$ is fed to Exp3 as a $[0, 1]$ -valued loss $\ell_t(I_t) = (2 - X_t)/4$, where $\mu_t(I_t) = A_t$. Since \mathcal{T}_2 is selected through independent coin tosses with fixed bias q , we can apply the Exp3 regret analysis (for losses chosen by a non-oblivious adversary) to bound the regret in \mathcal{T}_2 and show that

$$\sum_{t \in \mathcal{T}_2} \langle A_t^o - A_t, \theta_* \rangle = 4 \sum_{t \in \mathcal{T}_2} \left(\ell_t(I_t) - \ell_t(o) \right) = O(\sqrt{n \log n}) \quad (\text{C.1})$$

where we used $\mu_t(o) = A_t^o$ for all $t \in [T]$.

We have

$$\begin{aligned}R_{\mathcal{T}_2} &= \sum_{t \in \mathcal{T}_2} \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a - A_t, \theta_* \rangle \\ &= \sum_{t \in \mathcal{T}_2} \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a - A_t^o + A_t^o - A_t, \theta_* \rangle \\ &= \underbrace{\sum_{t \in \mathcal{T}_2} \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a - A_t^o, \theta_* \rangle}_{\text{Reg}_{\text{Img}}} + \underbrace{\sum_{t \in \mathcal{T}_2} \langle A_t^o - A_t, \theta_* \rangle}_{\text{Reg}_{\text{Exp3}}}.\end{aligned}$$

Reg_{Exp3} is bounded in (C.1), hence we only need to bound

$$\text{Reg}_{\text{Img}} = \sum_{t \in \mathcal{T}_2} \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a - A_t^o, \theta_* \rangle. \quad (\text{C.2})$$

Let

$$\tilde{\theta}_t^i = \operatorname{argmax}_{\theta \in \mathcal{C}_t^i} \max_{a \in \mathcal{A}_t} \langle a, \theta \rangle \quad \text{and} \quad A_t^* = \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a, \theta_* \rangle.$$

Recall \mathcal{E} is the event that $\theta_* \in \mathcal{C}_t$ for all $t \in [T]$. Assume \mathcal{E} holds. We obtain that for $t \in [T]$,

$$\begin{aligned}\langle \theta_*, A_t^* - A_t^o \rangle &\leq \langle \tilde{\theta}_t^o - \theta_*, A_t^o \rangle \quad (\text{because } \theta_* \in \mathcal{C}_t \subseteq \mathcal{C}_t^o) \\ &\leq \|\tilde{\theta}_t^o - \theta_*\|_{V_{t-1}} \|A_t^o\|_{V_{t-1}^{-1}} \quad (\text{Cauchy-Schwarz inequality}) \\ &\leq \sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}},\end{aligned} \quad (\text{C.3})$$

where the last inequality is because $\tilde{\theta}_t^o \in \mathcal{C}_t^o$. Then, assuming \mathcal{E} holds, we can bound Reg_{Img} as

$$\begin{aligned}\text{Reg}_{\text{Img}} &= \sum_{t \in \mathcal{T}_2} \langle \theta_*, A_t^* - A_t^o \rangle \\ &\leq \sum_{t \in \mathcal{T}_2} \min \left\{ 2, \sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}} \right\} \\ &\leq \sum_{t \in \mathcal{T}_2} \min \left\{ 2, \sqrt{2\gamma(1/T)} \|A_t^o\|_{V_{t-1}^{-1}} \right\} \quad (\text{due to } \alpha_o \leq 2\gamma(1/T)) \\ &\leq 2\sqrt{\gamma(1/T)} \sum_{t \in \mathcal{T}_2} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}} \right\},\end{aligned} \quad (\text{C.4})$$

where in the first inequality, we use the facts $\langle \theta_*, A_t^* - A_t^o \rangle \leq \sqrt{\alpha_o} \|A_t^o\|_{V_{t-1}^{-1}}$ and $\langle \theta_*, A_t^* - A_t^o \rangle \leq \langle \theta_*, A_t^* \rangle \leq 2$.

From Lemma D.3, we have

$$\sum_{t \in \mathcal{T}_2} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} \leq \sum_{t \in \mathcal{T}_2} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\}.$$

For each $s \geq 0$, let

$$\mathcal{T}_2^s = \left\{ t \in \mathcal{T}_2 : \det(V_{t-1}) \in \left[2^{ds}, 2^{d(s+1)} \right) \right\}$$

so that

$$\mathcal{T}_2 = \bigcup_{s \geq 0} \mathcal{T}_2^s.$$

Let $s' = \lceil \log_2 \det(V_T) / d \rceil$. We have

$$\begin{aligned} \mathbb{E}[\text{Reg}_{\text{Img}} \cdot \mathbb{1}[\mathcal{E}]] &\leq \mathbb{E} \left[2\sqrt{\gamma(1/T)} \sum_{t \in \mathcal{T}_2} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}} \right\} \right] && \text{(due to (C.4))} \\ &\leq 2\sqrt{2\alpha_o} \mathbb{E} \left[\sum_{t \in \mathcal{T}_2} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}} \right\} \right] && \text{(due to } \alpha_0 \leq 2\gamma(1/T) \text{)} \\ &\leq 2\sqrt{2\alpha_o T} \sqrt{\mathbb{E} \left[\sum_{t \in \mathcal{T}_2} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} \right]} && \text{(Cauchy-Schwarz inequality)} \\ &\leq 2\sqrt{2\alpha_o T} \sqrt{\sum_{s=1}^{s'} \mathbb{E} \left[\sum_{t \in \mathcal{T}_2^s} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} \right]} && \text{(where } d \ln s' \leq \log \det(V_T) \text{)} \\ &\leq 2\sqrt{2\alpha_o T} \sqrt{\sum_{s=1}^{s'} \mathbb{E} \left[\sum_{t \in \mathcal{T}_2^s} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} \right]} && \text{(because } \|A_t^n\|_{V_{t-1}^{-1}} \geq \|A_t^o\|_{V_{t-1}^{-1}} \text{)} \\ &\leq 2\sqrt{2\alpha_o T} \sqrt{\sum_{s=1}^{s'} (2d+1)/q} && \text{(due to Lemma D.4)} \\ &\leq 2\sqrt{2(2d+1)\alpha_o T s' / q}. \end{aligned}$$

To obtain the final results, we have

$$\det(V_T) = \prod_{i=1}^d \lambda_i \leq \left(\frac{1}{d} \text{trace}(V_T) \right)^d \leq \left(1 + \frac{TL^2}{d} \right)^d, \quad (\text{C.5})$$

where $\lambda_1, \dots, \lambda_d$ are the eigenvalues of V_T . Therefore, we have $s' \leq \lceil \log_2(1 + TL^2/d) \rceil$.

$$\begin{aligned} \mathbb{E}[\text{Reg}_{\text{Img}} \cdot \mathbb{1}[\mathcal{E}]] &\leq 2\sqrt{2(2d+1)\alpha_o T / q} \cdot \sqrt{s'} \\ &\leq 2\sqrt{(6d\alpha_o T / q) \lceil \log_2(1 + TL^2/d) \rceil}. \end{aligned}$$

By substituting the bounds on Reg_{Img} and Reg_{Exp3} into $R_{\mathcal{T}_2}$, we obtain

$$\begin{aligned} \mathbb{E}[R_{\mathcal{T}_2}] &= \mathbb{E}[\text{Reg}_{\text{Exp3}} \mathbb{1}[\mathcal{E}]] + \mathbb{E}[\text{Reg}_{\text{Img}}] + T \cdot \mathbb{P}(\mathcal{E}^c) \\ &\leq 2\sqrt{6d\alpha_o T / q} \sqrt{\lceil \log(1 + TL^2/d) \rceil} + O(\sqrt{nT \log n}), \end{aligned}$$

where the last inequality is because $\mathbb{P}(\mathcal{E}^c) \leq 1/T$ from Lemma 2.1 with our choice of $\delta = 1/\delta$.

Finally, we bound $R_{\mathcal{T}_1}$. Conditioned on event \mathcal{E} and using (D.3), $\forall t \in [T], \theta_* \in \mathcal{C}_t \subseteq \mathcal{C}_t^0 \subseteq \mathcal{C}_t^n$. Hence, we can obtain

$$\begin{aligned} \langle \theta_*, A_t^* - A_t^n \rangle &\leq \langle \tilde{\theta}_t^n - \theta_*, A_t^n \rangle && \text{(because } \theta_* \in \mathcal{C}_t \subseteq \mathcal{C}_t^n\text{)} \\ &\leq \|\tilde{\theta}_t^n - \theta_*\|_{V_{t-1}} \|A_t^n\|_{V_{t-1}^{-1}} && \text{(Cauchy-Schwarz inequality)} \\ &\leq \sqrt{\alpha_n} \|A_t^n\|_{V_{t-1}^{-1}} && \text{(because } \tilde{\theta}_t^n \in \mathcal{C}_t^n\text{.)} \end{aligned}$$

Hence, conditioned on event \mathcal{E} ,

$$\begin{aligned} R_{\mathcal{T}_1} &\leq \sum_{t \in \mathcal{T}_1} \min \{2, \langle \theta_*, A_t^* - A_t^n \rangle\} \\ &\leq \sum_{t \in \mathcal{T}_1} \min \left\{ 2, \sqrt{\alpha_n} \|A_t^n\|_{V_{t-1}^{-1}} \right\} && \text{(C.6)} \end{aligned}$$

$$\begin{aligned} &\leq 2\sqrt{\alpha_n |\mathcal{T}_1|} \sqrt{\sum_{t \in \mathcal{T}_1} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\}} && \text{(Cauchy-Schwarz inequality)} \\ &\leq 2\sqrt{\alpha_n |\mathcal{T}_1|} \sqrt{\sum_{t \in \mathcal{T}_1} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} + \sum_{t \in \mathcal{T}_2} \min \left\{ 1, \|A_t\|_{V_{t-1}^{-1}}^2 \right\}} \\ &= 2\sqrt{\alpha_n |\mathcal{T}_1|} \sqrt{\sum_{t \in [T]} \min \left\{ 1, \|A_t\|_{V_{t-1}^{-1}}^2 \right\}} \\ &\leq 2\sqrt{\alpha_n |\mathcal{T}_1|} \cdot \sqrt{2 \log(\det(V_T))}, && \text{(C.7)} \end{aligned}$$

where the last inequality is due to Lemma D.1. Therefore, the expected regret for $t \in \mathcal{T}_1$ is

$$\begin{aligned} \mathbb{E}[R_{\mathcal{T}_1}] &\leq \mathbb{E}[R_{\mathcal{T}_1} \mathbb{1}\{\mathcal{E}\}] + T \cdot \mathbb{P}(\mathcal{E}^c) \\ &\leq \mathbb{E} \left[2\sqrt{\alpha_n |\mathcal{T}_1|} \cdot \sqrt{2 \log(\det(V_T))} \right] + 1 \\ &\leq \mathbb{E} \left[2\sqrt{\alpha_n |\mathcal{T}_1|} \cdot \sqrt{2d \log \left(1 + \frac{TL^2}{d} \right)} \right] + 1 \\ &= 2\sqrt{\alpha_n} \mathbb{E} \left[\sqrt{|\mathcal{T}_1|} \right] \cdot \sqrt{2d \log \left(1 + \frac{TL^2}{d} \right)} + 1 \\ &\leq 2\sqrt{\alpha_n T q} \cdot \sqrt{2d \log \left(1 + \frac{TL^2}{d} \right)} + 1, && \text{(Jensen's inequality)} \end{aligned}$$

where the first inequality is due to (C.7) and $\mathbb{P}(\mathcal{E}^c) \leq 1/T$ from Lemma 2.1 and the last inequality is due to the fact that for any $t \in [T]$, with probability q , $t \in \mathcal{T}_1$. Finally, we obtain

$$\begin{aligned} \mathbb{E}[R_T] &= \mathbb{E}[R_{\mathcal{T}_1}] + \mathbb{E}[R_{\mathcal{T}_2}] \\ &\leq 2\sqrt{\alpha_n T q} \cdot \sqrt{2d \log \left(1 + \frac{TL^2}{d} \right)} + 2\sqrt{6d\alpha_o T/q} \sqrt{\lceil \log(1 + TL^2/d) \rceil} + O(\sqrt{nT \log n}). \end{aligned}$$

Note that $n = \Theta(\log d)$ and $\alpha_o = \Theta(S \log T)$. We obtain

$$\mathbb{E}[R_T] = O \left(\max \left\{ \sqrt{qd}, \sqrt{\frac{S}{q}} \right\} \cdot \sqrt{dT} \cdot \log T \right),$$

which completes the proof. \square

D Supporting lemmas

Lemma D.1 (Dani et al. [10]). *Let $A_1, A_2, \dots, A_T \in \mathbb{R}^d$ and $V_t = I + \sum_{s=1}^t A_s A_s^\top$ for all $t \in [T]$. Then*

$$\sum_{t=1}^T \min \left\{ 1, \|A_t\|_{V_{t-1}^{-1}}^2 \right\} \leq 2 \log \det(V_T). \quad (\text{D.1})$$

Moreover,

$$\min \left\{ 1, \|A_t\|_{V_{t-1}^{-1}}^2 \right\} \leq 2 \log \left(1 + \|A_t\|_{V_{t-1}^{-1}}^2 \right) = 2 \log \left(\frac{\det(V_t)}{\det(V_{t-1})} \right).$$

Lemma D.2. *For \mathcal{C}_t defined in (2.2), we have that $\mathcal{C}_t \subseteq \mathcal{C}_t^o$ for all $t \in [T]$.*

Proof. Recall

$$\mathcal{C}_{t+1} = \left\{ \theta : \|\theta\|_2^2 + \sum_{s=1}^t (\hat{X}_s - \langle \theta, A_s \rangle)^2 \leq \gamma(1/T) \right\}.$$

Note that

$$\begin{aligned} & \|\theta\|_2^2 + \sum_{s=1}^t (\hat{X}_s - \langle \theta, A_s \rangle)^2 - \|\hat{\theta}_t\|_2^2 - \sum_{s=1}^t (\hat{X}_s - \langle \hat{\theta}_t, A_s \rangle)^2 \\ &= \|\theta\|_2^2 - 2\theta^\top \left(\sum_{s=1}^t \hat{X}_s A_s \right) + \theta^\top \left(\sum_{s=1}^t A_s A_s^\top \right) \theta \\ & \quad - \|\hat{\theta}_t\|_2^2 + 2\hat{\theta}_t^\top \left(\sum_{s=1}^t \hat{X}_s A_s \right) - \hat{\theta}_t^\top \left(\sum_{s=1}^t A_s A_s^\top \right) \hat{\theta}_t \\ &= \|\theta\|_{V_t}^2 + 2(\hat{\theta}_t - \theta)^\top V_t \hat{\theta}_t - \|\hat{\theta}_t\|_{V_t}^2 \quad (V_t = I + \sum_{s=1}^t A_s A_s^\top, \hat{\theta}_t = V_t^{-1} \sum_{s=1}^t A_s \hat{X}_s) \\ &= \|\theta\|_{V_t}^2 - 2\theta^\top V_t \hat{\theta}_t + \|\hat{\theta}_t\|_{V_t}^2 = \|\theta - \hat{\theta}_t\|_{V_t}^2. \end{aligned}$$

Hence, we can express the ellipsoid as

$$\mathcal{C}_{t+1} = \left\{ \|\theta - \hat{\theta}_t\|_{V_t}^2 + \|\hat{\theta}_t\|_2^2 + \sum_{s=1}^t (\hat{X}_s - \langle \hat{\theta}_t, A_s \rangle)^2 \leq \gamma(1/T) \right\}. \quad (\text{D.2})$$

Therefore, for all $t \geq 0$,

$$\begin{aligned} \mathcal{C}_{t+1} &= \left\{ \theta : \|\theta - \hat{\theta}_t\|_{V_t}^2 + \|\hat{\theta}_t\|_2^2 + \sum_{s=1}^t (\hat{X}_s - \langle \hat{\theta}_t, A_s \rangle)^2 \leq \gamma(1/T) \right\} \quad (\text{due to (D.2)}) \\ &\subseteq \left\{ \theta : \|\theta - \hat{\theta}_t\|_{V_t}^2 \leq \gamma(1/T) \right\} \\ &\subseteq \left\{ \theta : \|\theta - \hat{\theta}_t\|_{V_t}^2 \leq \alpha_0 \right\} \quad (\text{by definition of } \alpha_0) \\ &= \mathcal{C}_{t+1}^o \end{aligned} \quad (\text{D.3})$$

concluding the proof. \square

Lemma D.3. *For any $1 \leq p \leq q \leq n$ and $t \in [T]$,*

$$\|A_t^p\|_{V_{t-1}^{-1}} \leq \|A_t^q\|_{V_{t-1}^{-1}}.$$

Proof. Note that

$$\begin{aligned} A_t^p &= \operatorname{argmax}_{a \in \mathcal{A}_t} \max_{\theta \in \mathcal{C}_t^p} \langle \theta, a \rangle = \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a, \hat{\theta}_{t-1} \rangle + \sqrt{\alpha_p} \|a\|_{V_{t-1}^{-1}}, \\ A_t^q &= \operatorname{argmax}_{a \in \mathcal{A}_t} \max_{\theta \in \mathcal{C}_t^q} \langle \theta, a \rangle = \operatorname{argmax}_{a \in \mathcal{A}_t} \langle a, \hat{\theta}_{t-1} \rangle + \sqrt{\alpha_q} \|a\|_{V_{t-1}^{-1}}. \end{aligned}$$

For contradiction, we assume $\|A_t^q\|_{V_{t-1}^{-1}} < \|A_t^p\|_{V_{t-1}^{-1}}$. Since $\|A_t^q\|_{V_{t-1}^{-1}} < \|A_t^p\|_{V_{t-1}^{-1}}$, $\|A_t^q\|_{V_{t-1}^{-1}} \neq \|A_t^p\|_{V_{t-1}^{-1}}$. Besides, according to the definition of A_t^p and A_t^q , we have

$$\langle A_t^q, \hat{\theta}_{t-1} \rangle + \sqrt{\alpha_p} \|A_t^q\|_{V_{t-1}^{-1}} \leq \langle A_t^p, \hat{\theta}_{t-1} \rangle + \sqrt{\alpha_p} \|A_t^p\|_{V_{t-1}^{-1}} \quad (\text{D.4})$$

$$< \langle A_t^p, \hat{\theta}_{t-1} \rangle + \sqrt{\alpha_q} \|A_t^p\|_{V_{t-1}^{-1}} \quad (\text{D.5})$$

$$\leq \langle A_t^q, \hat{\theta}_{t-1} \rangle + \sqrt{\alpha_q} \|A_t^q\|_{V_{t-1}^{-1}}, \quad (\text{D.6})$$

where the first inequality is due to the definition of A_t^p , the second inequality is due to $\sqrt{\alpha_p} < \sqrt{\alpha_q}$, and the last inequality is due to the definition of A_t^q . From the above results, we further have

$$\langle A_t^q - A_t^p, \hat{\theta}_{t-1} \rangle \leq \sqrt{\alpha_p} (\|A_t^p\|_{V_{t-1}^{-1}} - \|A_t^q\|_{V_{t-1}^{-1}}) \quad (\text{Due to (D.4)})$$

$$< \sqrt{\alpha_q} (\|A_t^p\|_{V_{t-1}^{-1}} - \|A_t^q\|_{V_{t-1}^{-1}})$$

(Due to assumption $\|A_t^q\|_{V_{t-1}^{-1}} < \|A_t^p\|_{V_{t-1}^{-1}}$ and $\alpha_q > \alpha_p$)

$$\leq \langle A_t^q - A_t^p, \hat{\theta}_{t-1} \rangle. \quad (\text{Due to (D.6).})$$

We obtain a contradiction. Therefore,

$$\|A_t^q\|_{V_{t-1}^{-1}} \geq \|A_t^p\|_{V_{t-1}^{-1}}. \quad (\text{D.7})$$

□

Lemma D.4.

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}_2^s} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq 2d/q + 1/q.$$

Proof. Let $\det(V_{t-1}) = (q_{t-1})^d$ and $\det(V_{t-1} + A_t^n (A_t^n)^\top) = (q_{t-1} + x_t)^d$. For each $s \geq 0$, let

$$\mathcal{T}^s = \left\{ t \in [T] : \det(V_{t-1}) \in [2^{sd}, 2^{d(s+1)}] \right\}.$$

Since $\mathcal{T}_2^s \subseteq \mathcal{T}^s$, to show

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}_2^s} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq 2d/q + 1/q,$$

we only need to prove

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}^s} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq 2d/q + 1/q.$$

We let $I_t = 1$ if the coin tosses in the t 's round is Head and 0 otherwise. We divide \mathcal{T}^s into two disjoint parts $\underline{\mathcal{T}^s}$ and $\overline{\mathcal{T}^s}$. Specifically,

- for $\underline{\mathcal{T}^s}$, it holds that for $t \in \underline{\mathcal{T}^s}$, $\det(V_{t-1} + A_t^n (A_t^n)^\top) \leq 2^{d(s+1)}$.
- for $\overline{\mathcal{T}^s}$, it holds that for $t \in \overline{\mathcal{T}^s}$, $\det(V_{t-1} + A_t^n (A_t^n)^\top) > 2^{d(s+1)}$.

From definition of $\underline{\mathcal{T}^s}$ and the fact that if $I_t = 1$, $V_t = V_{t-1} + A_t^n (A_t^n)^\top$, we have

$$\sum_{t \in \underline{\mathcal{T}^s}} x_t \cdot I_t \leq 2^s.$$

From Algorithm 2, with probability q , $I_t = 1$. Therefore, if we let $\{\mathcal{F}_t\}_{t \in [T]}$ be the natural filtration of $\{A_t, I_t, X_t\}_{t \in [T]}$, we have

$$2^s \geq \mathbb{E} \left[\sum_{t \in \underline{\mathcal{T}^s}} x_t \cdot I_t \right]$$

$$\begin{aligned}
&= \mathbb{E} \left[\sum_{t \in [T]} x_t \cdot I_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \right] \\
&= \sum_{t \in [T]} \mathbb{E}[x_t \cdot I_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\}] \\
&= \sum_{t \in [T]} \mathbb{E}[\mathbb{E}[x_t \cdot I_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} | \mathcal{F}_{t-1}]] \quad (\text{Law of total expectation}) \\
&= \sum_{t \in [T]} \mathbb{E}[x_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \cdot \mathbb{E}[I_t | \mathcal{F}_{t-1}]] \quad (x_t \text{ and } \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \text{ are } \mathcal{F}_{t-1}\text{-measurable}) \\
&= \sum_{t \in [T]} \mathbb{E}[x_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \cdot q] \quad (I_t \text{ is an independent coin-tossing}) \\
&= q \mathbb{E} \left[\sum_{t \in \underline{\mathcal{T}}^s} x_t \right]
\end{aligned}$$

and therefore $\mathbb{E} \left[\sum_{t \in \underline{\mathcal{T}}^s} x_t \right] \leq 2^s/q$. For $t \in \underline{\mathcal{T}}^s$, we further have

$$\begin{aligned}
\min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} &\leq 2 \log \left(\frac{\det(V_{t-1} + A_t^n (A_t^n)^\top)}{\det(V_{t-1})} \right) \quad (\text{due to Lemma D.1}) \\
&= 2d \log \left(1 + \frac{x_t}{q_{t-1}} \right) \\
&\leq 2d \log \left(1 + \frac{x_t}{2^s} \right) \quad (\text{since } t \in \mathcal{T}^s, q_{t-1} \geq 2^s) \\
&\leq \frac{2d \cdot x_t}{2^s} \quad (\text{due to } \log(1+x) \leq x \text{ for } x > 0.)
\end{aligned}$$

Therefore,

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}^s} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq \mathbb{E} \left[\sum_{t \in \mathcal{T}^s} \frac{2d \cdot x_t}{2^s} \right] = \frac{2d}{2^s} \mathbb{E} \left[\sum_{t \in \mathcal{T}^s} x_t \right] \leq 2d/q.$$

From definition of $\overline{\mathcal{T}}^s$, if $I_t = 1$ and $t \in \overline{\mathcal{T}}^s$, $\det(V_t) > 2^{d(s+1)}$. Then, for all $\tau > t$, $\tau \notin \mathcal{T}^s$. Therefore, there is at most one $t \in \overline{\mathcal{T}}^s$ with $I_t = 1$. We obtain

$$\mathbb{E} \left[\sum_{t \in \overline{\mathcal{T}}^s} \min \left\{ 1, \|A_t^n\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq \mathbb{E} \left[|\overline{\mathcal{T}}^s| \right] \leq 1/q,$$

where the last inequality is because with probability q , $I_t = 1$. By combining the bounds for $t \in \overline{\mathcal{T}}^s$ and $t \in \underline{\mathcal{T}}^s$ together, lemma follows. \square

Lemma D.5. Let $Q = \sum_{s \geq o} q_s$. Then,

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}^s} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq 2d/Q + 1/Q.$$

Proof. Let $\det(V_{t-1}) = (q_{t-1})^d$. Let x_t satisfies $\det(V_{t-1} + A_t^o (A_t^o)^\top) = (q_{t-1} + x_t)^d$. We let $I_t = \mathbb{1}\{I_t \geq o\}$. We divide \mathcal{T}^s into two disjoint parts $\underline{\mathcal{T}}^s$ and $\overline{\mathcal{T}}^s$. Specifically,

- for $\underline{\mathcal{T}}^s$, it holds that for $t \in \underline{\mathcal{T}}^s$, $\det(V_{t-1} + A_t^o (A_t^o)^\top) \leq 2^{d(s+1)}$.
- for $\overline{\mathcal{T}}^s$, it holds that for $t \in \overline{\mathcal{T}}^s$, $\det(V_{t-1} + A_t^o (A_t^o)^\top) > 2^{d(s+1)}$.

Note that for $I_t \geq o$,

$$\begin{aligned}
& \log(\det(V_t)) - \log(\det(V_{t-1} + A_t^o (A_t^o)^\top)) \\
&= \log(\det(V_t)) - \left(\log(\det(V_{t-1} + A_t^o (A_t^o)^\top)) - \log(V_{t-1}) \right) \\
&\quad \text{(due to Lemma D.1)} \\
&= \log(1 + \|A_t^{I_t}\|_{V_{t-1}^{-1}}) - \log(1 + \|A_t^o\|_{V_{t-1}^{-1}}) \\
&\geq 0. \quad \text{(due to Lemma D.3)}
\end{aligned}$$

Therefore, if we let $\det(V_{t-1} + A_t^{I_t} (A_t^{I_t})^\top) = (q_{t-1} + x'_t)^d$ and $I_t \geq o$, then $x'_t \geq x_t$. Hence,

$$\sum_{t \in \underline{\mathcal{T}}^s} x_t \cdot I_t \leq 2^s.$$

Note that $\mathbb{P}(I_t \geq o) = \sum_{s \geq o} q_s = Q$. Let $\{\mathcal{F}_t\}_{t \in [T]}$ be the natural filtration of $\{A_t, I_t, X_t\}_{t \in [T]}$. We have

$$\begin{aligned}
2^s &\geq \mathbb{E} \left[\sum_{t \in \underline{\mathcal{T}}^s} x_t \cdot I_t \right] \\
&= \mathbb{E} \left[\sum_{t \in [T]} x_t \cdot I_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \right] \\
&= \sum_{t \in [T]} \mathbb{E}[x_t \cdot I_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\}] \\
&= \sum_{t \in [T]} \mathbb{E}[\mathbb{E}[x_t \cdot I_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} | \mathcal{F}_{t-1}]] \quad \text{(Law of total expectation)} \\
&= \sum_{t \in [T]} \mathbb{E}[x_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \cdot \mathbb{E}[I_t | \mathcal{F}_{t-1}]] \quad (x_t \text{ and } \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \text{ are } \mathcal{F}_{t-1}\text{-measurable}) \\
&= \sum_{t \in [T]} \mathbb{E}[x_t \cdot \mathbb{1}\{t \in \underline{\mathcal{T}}^s\} \cdot Q] \quad (I_t \text{ is an independent coin-tossing}) \\
&= Q \mathbb{E} \left[\sum_{t \in \underline{\mathcal{T}}^s} x_t \right]
\end{aligned}$$

and therefore

$$\mathbb{E} \left[\sum_{t \in \underline{\mathcal{T}}^s} x_t \right] \leq 2^s / Q.$$

For $t \in \underline{\mathcal{T}}^s$, we further have

$$\begin{aligned}
\min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} &\leq \min \left\{ 1, \|A_t^{I_t}\|_{V_{t-1}^{-1}}^2 \right\} \\
&\leq 2 \log \left(\frac{\det(V_{t-1} + A_t^{I_t} (A_t^{I_t})^\top)}{\det(V_{t-1})} \right) \quad \text{(due to Lemma D.1)} \\
&= 2d \log \left(1 + \frac{x_t}{q_{t-1}} \right) \\
&\leq 2d \log \left(1 + \frac{x_t}{2^s} \right) \quad (\text{since } t \in \mathcal{T}^s, q_{t-1} \geq 2^s) \\
&\leq \frac{2d \cdot x_t}{2^s} \quad (\text{due to } \log(1 + x) \leq x \text{ for } x > 0.)
\end{aligned}$$

Therefore,

$$\mathbb{E} \left[\sum_{t \in \mathcal{T}^s} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq \mathbb{E} \left[\sum_{t \in \mathcal{T}^s} \frac{2d \cdot x_t}{2^s} \right] = \frac{2d}{2^s} \mathbb{E} \left[\sum_{t \in \mathcal{T}^s} x_t \right] \leq 2d/Q.$$

From definition of $\overline{\mathcal{T}^s}$, if $I_t = 1$ and $t \in \overline{\mathcal{T}^s}$, then $I_t \geq o$ and

$$\det(V_t) = \det(V_{n-1} + A_t^{I_t} (A_t^{I_t})^\top) > \det(V_{n-1} + A_t^o (A_t^o)^\top) > (s+1)^d.$$

Then, for all $\tau > t$, $\tau \notin \mathcal{T}^s$. Therefore, there is at most one $t \in \overline{\mathcal{T}^s}$ with $I_t = 1$. We obtain

$$\mathbb{E} \left[\sum_{t \in \overline{\mathcal{T}^s}} \min \left\{ 1, \|A_t^o\|_{V_{t-1}^{-1}}^2 \right\} \right] \leq \mathbb{E} \left[|\overline{\mathcal{T}^s}| \right] \leq 1/Q,$$

where the last inequality is because with probability $\sum_{s \geq o} q_s = Q$, $I_t = 1$. By combining the bounds for $t \in \overline{\mathcal{T}^s}$ and $t \in \mathcal{T}^s$ together, lemma follows. \square

E Experimental details

The code used in the experiments can be found in the following repository: https://github.com/jajajang/sparsity_agnostic_model_selection.

E.1 Settings common to all algorithms

- Arm set: $\mathcal{A}_t = \mathcal{A}$ for all $t \in [T]$. $\mathcal{A} \subset \mathbb{S}^{d-1}$ (the unit sphere in \mathbb{R}^d) is a set of d -dimensional vectors drawn independently and uniformly at random from \mathbb{S}^{d-1} , with $d = 16$ and $|\mathcal{A}| = 30$.
- θ_* is an S -sparse ($S = 1, 2, 4, 8, 16$) vector generated as follows: before the game starts, draw $(\theta_*)_1, \dots, (\theta_*)_S \sim \mathbb{S}^{S-1}$, and $(\theta_*)_k = 0$ for all $k > S$.
- The noise on rewards: $\{\varepsilon_t\}_{t \in [T]}$ are i.i.d. with $\xi_t \sim \text{Unif}([-1, 1])$.
- Number of iterations: $T = 10^4$.
- Number of models: $n = 6$.
- Radius of confidence sets: $\alpha_0 = 0$, and $\alpha_i = 2^i \log t$ for $i = 1, \dots, 5$.
- Prior distribution $\{q_s\}_{s \in [6]}$
 - For `_Unif`, $\{q_s\}_{s \in [6]} = (\frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6})$
 - For `_Theory`, $\{q_s\}_{s \in [6]} = (\frac{C}{2}, \frac{C}{4}, \frac{C}{8}, \frac{C}{16}, \frac{C}{32}, \frac{C}{64})$ where $C = \frac{63}{64}$ is a normalizing constant.
- Each plot is the result of 20 repetitions for each method. The shade represents the 1-standard deviation bound.
- Hardware: Lenovo Thinkpad P16s Gen 2 Laptop - Type 21HL
 - CPU: 13th Gen Intel(R) Core(TM) i7-1360P 2.20 GHz
 - RAM: 32GB
- Computation time: total 1338.38 seconds.

E.2 AdaLinUCB details

- Since we empirically observed that Exp3 provided enough exploration, we aggressively set the forced exploration parameter q to zero.
- The learning rate of Exp3 was set to $\eta_t = 2\sqrt{\frac{\log n}{nt}}$, see [5].
- Given the prior distribution $\{q_s\}_{s \in [6]}$, we set P_t as follows:

$$P_{t,s} = \frac{q_s \exp(\eta_t S_{t,s})}{\sum_{j=1}^n q_j \exp(\eta_t S_{t,j})}$$

E.3 OFUL details

- We used the log-determinant form of the confidence set based on Abbasi-Yadkori et al. [1, Theorem 2], which gives the choice

$$\sqrt{\gamma_t} = \sqrt{2 \log T + \log \det(V_t)} + 1$$

for the parameter γ_t in (1.1) when $\lambda = 1$ and $\delta = 1/T$.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: **[Yes]**

Justification: The main claims presented in the abstract and introduction accurately represent the contributions and scope of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: **[Yes]**

Justification: We have added a paragraph in the conclusions section discussing the limitations of our work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We provide the full set of assumptions and complete proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We provide all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: We include our code in our supplementary material and will make the full code public if this work gets accepted.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: We provide all the details of our experimental setting.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: We reported 1-sigma error bars in our plots.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: **[Yes]**

Justification: We state all information on the computer resources in Appendix E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: **[Yes]**

Justification: Our research conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: **[NA]**

Justification: There is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Our paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.