

---

# Syndicated Bandits: A Framework for Auto Tuning Hyper-parameters in Contextual Bandit Algorithms

---

**Qin Ding**

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

**Yue Kang**

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

**Yi-Wei Liu**

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

**Thomas C. M. Lee**

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

**Cho-Jui Hsieh**

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

**James Sharpnack**

Amazon\*  
Berkeley, CA  
jsharpna@gmail.com

## Abstract

The stochastic contextual bandit problem, which models the trade-off between exploration and exploitation, has many real applications, including recommender systems, online advertising and clinical trials. As many other machine learning algorithms, contextual bandit algorithms often have one or more hyper-parameters. As an example, in most optimal stochastic contextual bandit algorithms, there is an unknown exploration parameter which controls the trade-off between exploration and exploitation. A proper choice of the hyper-parameters is essential for contextual bandit algorithms to perform well. However, it is infeasible to use offline tuning methods to select hyper-parameters in contextual bandit environment since there is no pre-collected dataset and the decisions have to be made in real time. To tackle this problem, we first propose a two-layer bandit structure for auto tuning the exploration parameter and further generalize it to the Syndicated Bandits framework which can learn multiple hyper-parameters dynamically in contextual bandit environment. We derive the regret bounds of our proposed Syndicated Bandits framework and show it can avoid its regret dependent exponentially in the number of hyper-parameters to be tuned. Moreover, it achieves optimal regret bounds under certain scenarios. Syndicated Bandits framework is general enough to handle the tuning tasks in many popular contextual bandit algorithms, such as LinUCB, LinTS, UCB-GLM, etc. Experiments on both synthetic and real datasets validate the effectiveness of our proposed framework.

## 1 Introduction

The stochastic contextual bandit problem models the well-known exploration-exploitation dilemma in a repeated game between a player and an environment. At each round, the player sequentially interacts with the environment by pulling an arm from a pool of  $K$  arms, where every arm is associated with a

---

\*Work done prior to joining Amazon.

$d$ -dimensional contextual feature vector. Only the stochastic reward corresponding to the pulled arm is revealed to the player. The goal of the player is to maximize the cumulative reward or minimize the cumulative regret. Due to the partial feedback setting, the player has to balance between exploitation — pulling the arm that has the best estimated reward so far — and exploration — exploring whether there are uncertain arms that may be better than the current estimated best arm.

With substantial applications in recommender systems [7], online advertising [20], clinical trials [25], etc., bandit algorithms have been extensively studied during the past few decades. In general, there are two exploration techniques, Upper Confidence Bound (UCB) [6, 17, 18] and Thompson Sampling (TS) [4, 5] algorithms. The UCB algorithm addresses the dilemma optimistically by pulling the arm that has the biggest upper confidence bound. The TS algorithm usually assumes that the relationship between the contextual features and rewards follows a prior model and it uses new observations at each round to estimate the posterior model. The player then pulls the arm that has the best estimated reward based on the posterior model. In general, the contextual bandit problems have regret lower bounded by  $O(\sqrt{T})$  [13, 17], where  $T$  is the total number of rounds. Both UCB and TS algorithms have been shown to achieve optimal regret bounds in (generalized) linear bandit problems [17, 18], kernelized bandit problems [2] and even in the contextual bandit problem with more complicated models such as neural networks [28].

Despite the popularity of the contextual bandit problem, there are some practical issues that prevent it from being used widely in practice. In both UCB and TS, there are hyper-parameters that are unknown to the player. One of the most important hyper-parameters is the exploration parameter, which controls the trade-off between exploration and exploitation. A good choice of the exploration parameter is essential for the algorithm to perform well and for the theory to hold. Another commonly seen hyper-parameter is the regularization parameter  $\lambda$  in ridge regression or generalized linear model, which is used to model the relationship between features and rewards in (generalized) linear bandits. In contextual bandit problems with complex models such as neural network, the recently proposed NeuralUCB [23] algorithm has far more than just two hyper-parameters. NeuralUCB also needs to select the network width, network depth, step size for gradient descent to solve the neural networks and gradient descent steps, etc. Due to the nature of the bandit environment, where the decisions have to be made in real time, it is inherently difficult to tune the hyper-parameters by the traditional offline tuning methods, such as cross validation, since when you have decided to use a parameter in partial datasets and make a decision based on this, the regret incurred by this decision will never be reversible in contextual bandit environment. In many prominent bandit works [17, 14, 28, 10], the experiments are conducted by running a grid search on the possible choices of parameters and only the best result is reported. Although the performance of the best grid search results is of academic interest, it is not possible to do grid search in practice. In some other works [15], exploration parameter is set as a sufficient, theoretically derived, and often unknown value, but this value may be too conservative and may not achieve good performances in practice, which can be seen from the experiments in Table 1.

In this work, we first propose a two-layer bandit structure that can automatically tune the exploration parameter dynamically from the observed data. The two-layer bandit structure has its similarities to the bandit-over-bandit (BOB) algorithm [11] proposed for the non-stationary stochastic bandit problems, where it uses BOB idea to successfully adapt its sliding-window sizes by restarting the algorithms in epochs. Motivated by the two-layer bandit structure we propose in Section 4, we generalize it to the “Syndicated Bandits” framework where there could be multiple hyper-parameters to be tuned in the contextual bandit algorithm. We provide theoretical guarantees for our framework and show that our proposed auto tuning method in general has regret upper bound  $\tilde{O}(T^{2/3}) + \tilde{O}(\sum_{l=1}^L \sqrt{n_l T})$ . Here  $L$  is the total number of hyper-parameters to be tuned and  $n_l$  is the number of candidates in the tuning set of the  $l$ -th hyper-parameter. When the unknown theoretical exploration parameter is no bigger than any elements in the tuning set, our proposed framework has optimal regret upper bound  $\tilde{O}(\sqrt{T}) + \tilde{O}(\sum_{l=1}^L \sqrt{n_l T})$  for UCB-based algorithms. Our framework is general enough to handle tuning tasks in many contextual bandit algorithms as long as the arms to be pulled at round  $t$  follows a fixed distribution given the hyper-parameters to be used at this round and the past information. This includes many popular contextual bandit algorithms such as Linear UCB (LinUCB) [17, 11], Linear TS (LinTS) [5, 10], UCB-GLM [18], etc. Our proposed Syndicated Bandits framework is the first work that considers tuning multiple parameters dynamically from observations in the contextual bandit problems with theoretical guarantees. We provide a regret bound that avoids the exponential dependency on the total number of hyper-parameters to be tuned. This is one of the main contributions of our proposed work. In Section 6, we show by experiments that our proposed framework improves

over existing works, as well as the bandit algorithms that use the unknown theoretically derived exploration parameter.

## 2 Related work

There is a rich line of works on multi-armed bandit (MAB) and stochastic contextual bandit algorithms, including (generalized) linear bandits, kernelized bandits and neural bandits, etc. Most of them follow the UCB and TS exploration techniques. We refer the readers to [17, 18, 4, 5, 10, 12, 28] for the seminal works regarding the bandit problems. There are many previous works that utilize algorithms in the stochastic MAB [23] setting to solve the hyper-parameter optimization problem [21, 19]. There are also some online hyper-parameter tuning works such as [24], however, those mainly focuses on reducing the training cost for tuning parameters of neural networks online and they are not considering minimizing the cumulative regret in contextual bandit problems. In the following, we will only pay attention to related works on the tuning tasks in stochastic contextual bandits.

[22] proposed a meta-learning method for learning exploration parameters in contextual bandit problems. It learns a good exploration strategy in synthetic datasets and applies it to the real contextual bandit problems by an imitation study. The meta-learning algorithm is compared with seven baseline contextual bandit algorithms and achieves good empirical results. We note that this algorithm cannot learn the exploration parameters adaptively from observations in the contextual bandit environment. In [9], the authors first proposed OPLINUCB and DOPLINUCB algorithms to learn exploration parameters dynamically. OPLINUCB treats the possible choices of hyper-parameters as arms and uses a standard MAB TS algorithm to choose parameters. It then uses the chosen parameter in the contextual bandit algorithm. However, this method does not have theoretical guarantee in general, since the MAB TS only works when the rewards of the candidate hyper-parameters in the tuning set stay stationary over time. For hyper-parameter selections in contextual bandit problems, the best exploration parameter does not stay the same all the time. This is because in later rounds, when the learning is sophisticated, less exploration is better. However, in the beginning, more exploration is preferred due to the uncertainty. This non-stationary nature in tuning hyper-parameters makes the performance of OPLINUCB unstable in practice. DOPLINUCB is a similar tuning method as OPLINUCB, except that it uses the CTree algorithm to select hyper-parameters at each round. It is shown in [9] that DOPLINUCB does not outperform OPLINUCB in stationary contetxual bandit environments, where the reward-feature model does not change over time.

Another close line of literature is on model selections in bandit algorithms. [16] tackles the feature selection problem in bandit algorithms and achieve  $O(T^{2/3}d_*^{1/3})$  where  $d_*$  is the total number of optimal features. [3] uses the coralling idea to create a master algorithm to choose the best bandit model from a set of  $M$  base models. Hyper-parameter tuning problem can be formulated as a model selection problem in [3], where we can treat bandit algorithms with different hyper-parameters as the base models. The theoretical regret bound of the coralling idea [3] is  $O(\sqrt{MT} + MR_{\max})$ , where  $M$  is the total number of base models and  $R_{\max}$  is the maximum regret of  $M$  base models if they were to run alone. This means that the regret bound will be exponentially dependent on the total number of hyper-parameters to be tuned. In addition, if there is one hyper-parameter in the tuning set that gives linear regret of the algorithm, then  $R_{\max}$  is linear in  $T$  which makes the coralling idea have linear regret in worst case. Our algorithm is also much more efficient than the coralling idea when  $M$  is big. The coralling idea requires updating all  $M$  base models/ algorithms at each round. However, our algorithm only needs to update the selected model/ bandit algorithm with selected hyper-parameter at each round. When the time complexity of updating the model/ algorithm is big, the coralling idea is expensive. For example, if we tune configurations for UCB-GLM, the coralling idea needs  $O(MT^2d)$  time, while the time complexity of our algorithm is only  $O(MT + T^2d)$ .

We address here that none of the previous works can tune multiple parameters dynamically from observations. Although OPLINUCB [9] and the coralling idea [3] can treat all the hyper-parameters as a single parameter and set the tuning set as all the possible combinations of hyper-parameters. This will lead to exponential number of configurations which may not be efficient in both computation and theoretical regret bounds. Our proposed Syndicated framework avoids the exponential regret bound.

**Notations:** For a vector  $x \in \mathbb{R}^d$ , we use  $\|x\|$  to denote its  $l_2$  norm and  $\|x\|_A := \sqrt{x^T A x}$  for a positive-definite matrix  $A \in \mathbb{R}^{d \times d}$ . Finally, we denote  $[n] := \{1, 2, \dots, n\}$ .

### 3 Preliminaries

We study the hyper-parameter selection tasks in a stochastic contextual bandit problem with  $K$  arms, where  $K$  can be an infinite number. Assume there are in total  $T$  rounds, at each round  $t \in [T]$ , the player is given  $K$  arms, represented by a set of feature vectors  $\mathcal{A}_t = \{x_{t,a} | a \in [K]\} \subset \mathbb{R}^d$ ,  $\mathcal{A}_t$  is drawn IID from an unknown distribution with  $\|x_{t,a}\| \leq 1$  for all  $t \in [T]$  and  $a \in [K]$ , where  $x_{t,a}$  is a  $d$ -dimensional feature vector that contains the information of arm  $a$  at round  $t$ . The player makes a decision by pulling an arm  $a_t \in [K]$  based on  $\mathcal{A}_t$  and past observations. We make a common regularity assumption as in [14, 18], i.e. there exists a constant  $\sigma_0 > 0$  such that  $\lambda_{\min} \left( \mathbb{E} \left[ \frac{1}{k} \sum_{a=1}^k x_{t,a} x_{t,a}^\top \right] \right) > \sigma_0$ . The player can only observe the rewards of the pulled arms. Denote  $X_t := x_{t,a_t}$  as the feature vector of the pulled arm at round  $t$  and  $Y_t$  the corresponding reward. We assume the expected rewards and features follow a model  $\mathbb{E}[Y_t | X_t] = \mu(X_t^\top \theta^*)$ , where  $\mu(\cdot)$  is a known model function and  $\theta^*$  is the true but unknown model parameter. When  $\mu(x) = x$ , this becomes the well-studied linear bandits problem. When  $\mu(\cdot)$  is a generalized linear model or a neural network, this becomes the generalized linear bandits (GLB) and neural bandits respectively.

Without loss of generality, we assume that there exists a positive constant  $S$  such that  $\|\theta^*\| \leq S$ . We also assume the mean rewards  $\mu(x_{t,a}^\top \theta^*) \in [0, 1]$  and observed rewards  $Y_t \in [0, 1]$ . This is a non-critical assumption, which can be easily relaxed to any bounded interval. If  $\mathcal{F}_t = \sigma(\{a_s, \mathcal{A}_s, Y_s\}_{s=1}^t \cup \mathcal{A}_{t+1})$  is the information up to round  $t$ , we assume the observed rewards follow a sub-Gaussian distribution with parameter  $\sigma^2$ , i.e.,  $Y_t = \mu(X_t^\top \theta^*) + \epsilon_t$ , where  $\epsilon_t$  are independent random noises that satisfy  $\mathbb{E}[e^{b\epsilon_t} | \mathcal{F}_{t-1}] \leq \frac{b^2 \sigma^2}{2}$  for all  $t$  and  $b \in \mathbb{R}$ . Denote  $a_t^* = \operatorname{argmax}_{a \in [K]} \mu(X_t^\top \theta^*)$  as the optimal arm at round  $t$  and  $x_{t,*}$  as its corresponding feature, the goal is to minimize the cumulative regret over  $T$  rounds defined in the following equation.

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

For linear bandits where  $\mu(x) = x$ , classic bandit algorithms such as LinUCB [1, 17] and LinTS [2] compute an estimate of the model parameter  $\hat{\theta}_t$  using ridge regression with regularization parameter  $\lambda > 0$ , i.e.,  $\hat{\theta}_t = V_t^{-1} \sum_{s=1}^{t-1} X_s Y_s$ , where  $V_t = \lambda I_d + \sum_{s=1}^{t-1} X_s X_s^\top$ . Shown by [1], with probability at least  $1 - \delta$ , the true model parameter  $\theta^*$  is contained in the following confidence set,

$$\mathcal{C}_t = \left\{ \theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_t\|_{V_t} \leq \alpha(t) \right\}, \quad (2)$$

where

$$\alpha(t) = \sigma \sqrt{d \log \left( \frac{1+t/\lambda}{\delta} \right)} + S\sqrt{\lambda}. \quad (3)$$

To balance the trade-off between exploration and exploitation, there are in general two techniques. For example, in linear bandits, LinUCB explores optimistically by pulling the arm with the maximum upper confidence bound, while LinTS adds randomization by drawing a sample model from the posterior distribution and pulls an arm based on it.

$$a_t = \operatorname{argmax}_a x_{t,a}^\top \hat{\theta}_t + \alpha(t) \|x_{t,a}\|_{V_t^{-1}}, \quad (\text{LinUCB})$$

$$\theta_t^{\text{TS}} \sim N(\hat{\theta}_t, \alpha(t)^2 V_t^{-1}) \quad \text{and} \quad a_t = \operatorname{argmax}_a x_{t,a}^\top \theta_t^{\text{TS}}. \quad (\text{LinTS})$$

In the following, we call  $\alpha(t)$  the exploration parameter. As suggested by the theories in [1, 17], a conservative choice of the exploration parameter is to follow Equation 3. However, in Equation 3 the upper bound of the  $l_2$  norm of the model parameter  $S$  and the sub-Gaussian parameter  $\sigma$  are unknown to the player, which makes it difficult to track theoretical choices of the exploration parameter.

In Table 1, we show the cumulative regret of LinUCB [1, 17] and LinTS [5] in a simulation study with  $d = 5$ ,  $T = 10000$  and  $K = 100$ . Rewards are simulated from  $N(x_{t,a}^\top \theta^*, 0.5)$ . The model parameter  $\theta^*$  and feature vectors  $x_{t,a}$  are drawn from  $\text{Uniform}(-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}})$ . Two scenarios are considered in this table. In the first scenario, the feature of each arm keeps the same over  $T$  rounds. While in the second scenario, the features are re-simulated from  $\text{Uniform}(-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}})$  at different rounds. We

run a grid search of the exploration parameter in  $\{0, 0.5, 1, \dots, 10\}$  and report the best grid search result, as well as the results using the theoretical exploration parameter given by Equation 3 (last column in Table 1). As we shall see in Table 1, the best exploration parameter is not the same for different scenarios. Therefore, which exploration parameter to use is an instance-dependent problem and the best exploration parameter should always be chosen dynamically based on the observations. Meanwhile, theoretical exploration parameters do not always give the best performances from Table 1. On the other hand, in many other works where the model of contextual bandit problem is more complex, such as the generalized linear bandit [14], neural bandit [28], there may be many more hyper-parameters than just  $\alpha(t)$ .

Table 1: Averaged cumulative regret (cum. reg.) and standard deviation (std) of the cumulative regret based on 5 repeated experiments. “Fixed” means that the feature vectors of each arm are fixed over time. “Changing” means that the features are re-simulated from the same distribution at each round.

Algorithm	Feature type	Best $\alpha(t)$	Cum. reg. (std)	Theoretical cum. reg. (std)
LinUCB	Fixed	4.0	357.21 (188.72)	364.99 (151.54)
	Changing	1.5	312.69 (42.53)	582.59 (523.60)
LinTS	Fixed	1.5	336.44 (137.01)	576.83 (110.48)
	Changing	3.5	352.79 (109.84)	488.99 (141.34)

## 4 A two-layer bandit structure for tuning exploration parameters

In the previous section, we have discussed that the best hyper-parameters should be instant-dependent. In this section, we propose a two-layer bandit structure to automatically learn the best hyper-parameter from data at each round. We will take learning the best exploration parameter as an example. However, we want to emphasize that this structure can also be applied to learn other single hyper-parameter.

We randomly select arms for the first  $T_1$  rounds to warm up the algorithm. For all rounds later, in this two-layer bandit structure, the top layer follows an adversarial MAB policy, namely, the EXP3 algorithm [7]. Assume  $J$  is the tuning set of all the possible exploration parameters. At each round  $t > T_1$ , the top layer will select a candidate exploration parameter  $\alpha_{i_t} \in J$ , where  $\alpha_i$  is the  $i$ -th element in the set  $J$  and  $i_t$  is the selected index at round  $t$ . The bottom layer runs the contextual bandit algorithm based on the selected exploration parameter  $\alpha_{i_t}$ . Details are listed in Algorithm 1.

### 4.1 Regret analysis

Given all the past information  $\mathcal{F}_{t-1}$ , denote  $a_t(\alpha_j | \mathcal{F}_{t-1})$  as the pulled arm when the exploration parameter is  $\alpha_j$  at round  $t$ . Denote  $X_t(\alpha_j | \mathcal{F}_{t-1}) = x_{t,a_t(\alpha_j | \mathcal{F}_{t-1})}$  as the corresponding feature vector under  $\mathcal{F}_{t-1}$ . Note that in our algorithm,  $X_t := X_t(\alpha_{i_t} | \mathcal{F}_{t-1})$  when  $t > T_1$ . To analyze the cumulative regret, we first decompose the regret defined in Equation 1 into three parts:

$$\begin{aligned} \mathbb{E}[R(T)] &= \mathbb{E} \left[ \sum_{t=1}^T (\mu(x_{t,*}^T \theta) - \mu(X_t^T \theta)) \right] = \underbrace{\mathbb{E} \left[ \sum_{t=T_1+1}^T (\mu(x_{t,*}^T \theta) - \mu(X_t(\alpha^* | \mathcal{F}_{t-1})^T \theta)) \right]}_{\text{Quantity (A)}} \\ &+ \underbrace{\mathbb{E} \left[ \sum_{t=T_1+1}^T (\mu(X_t(\alpha^* | \mathcal{F}_{t-1})^T \theta) - \mu(X_t(\alpha_{i_t} | \mathcal{F}_{t-1})^T \theta)) \right]}_{\text{Quantity (B)}} + \underbrace{\mathbb{E} \left[ \sum_{t=1}^{T_1} (\mu(x_{t,*}^T \theta) - \mu(X_t^T \theta)) \right]}_{\text{Quantity (C)}}, \end{aligned}$$

where  $\mu(\cdot)$  is the reward-feature model function and  $\alpha^* \in J$  is some arbitrary candidate exploration parameter in  $J$ . Quantity (A) is the regret of the contextual bandit algorithm that runs with the same hyper-parameter  $\alpha^*$  under the past history  $\mathcal{F}_{t-1}$  generated from our tuning strategy every round. Quantity (B) is the extra regret paid to tune the hyper-parameter. Quantity (C) is the regret paid for random exploration in warm-up phases and is controlled by the scale of  $O(T_1)$ . We show by Lemma 1 and Theorem 1 below that our auto tuning method in Algorithm 1 does not cost too much in selecting parameters in most scenarios under mild conditions.

---

**Algorithm 1** A Two-layer Auto Tuning Algorithm

---

**Input:** time horizon  $T$ , warm-up length  $T_1$ , candidate hyper-parameter set  $J = \{\alpha_i\}_{i=1}^n$ .

- 1: Randomly choose  $a_t \in [K]$  and record  $X_t, Y_t$  for  $t \in [T_1]$ .
- 2: Initialize exponential weights  $w_j(T_1 + 1) = 1$  for  $j = 1, \dots, n$ .
- 3: Initialize the exploration parameter for EXP3 as  $\beta = \min \left\{ 1, \sqrt{\frac{n \log n}{(e-1)T}} \right\}$ .
- 4: **for**  $t = (T_1 + 1)$  **to**  $T$  **do**
- 5:     Update probability distribution for pulling candidates in  $J$

$$p_j(t) = \frac{\beta}{n} + (1 - \beta) \frac{w_j(t)}{\sum_{i=1}^n w_i(t)}$$

- 6:      $i_t \leftarrow j \in [n]$  with probability  $p_j(t)$ .
- 7:     Run the contextual bandit algorithm with hyper-parameter  $\alpha(t) = \alpha_{i_t}$  to pull an arm. For example, pull arms according to the following equations

$$a_t = \underset{a=1, \dots, K}{\operatorname{argmax}} x_{t,a}^T \hat{\theta}_t + \alpha_{i_t} \|x_{t,a}\|_{V_t^{-1}} \quad (\text{LinUCB})$$

$$\theta_t^{\text{TS}} \sim N(\hat{\theta}_t, \alpha_{i_t}^2 V_t^{-1}) \quad \text{and} \quad a_t = \underset{a}{\operatorname{argmax}} x_{t,a}^T \theta_t^{\text{TS}}. \quad (\text{LinTS})$$

- 8:     Observe reward  $Y_t$  and update the components in the contextual bandit algorithm.
- 9:     Update EXP3 components:  $\hat{y}_t(j) \leftarrow 0$  if  $j \neq i_t$ ,  $\hat{y}_t(j) \leftarrow Y_t/p_j(t)$  if  $j = i_t$ , and

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

---

10: **end for**

---

Since the arms pulled by the contextual bandit layer also affect the update of the EXP3 layer in Algorithm 1, the result of EXP3 algorithm is not directly applicable to bounding Quantity (B). We modify the proof techniques in [7] and present the proof details in Appendix.

**Lemma 1.** *Assume given the past information  $\mathcal{F}_{t-1}$  and the hyper-parameter to be used by the contextual bandit algorithm at round  $t$ , the arm to be pulled follows a fixed distribution. For a random sequence of hyper-parameters  $\{\alpha_{i_1}, \dots, \alpha_{i_T}\}$  selected by the EXP3 layer in Algorithm 1 and arm  $a_t(\alpha_{i_t})$  is pulled in the contextual bandit layer at round  $t$ , we have*

$$\max_{\alpha \in J} \mathbb{E} \left[ \sum_{t=1}^T \mu \left( X_t(\alpha | \mathcal{F}_{t-1})^T \theta \right) \right] - \mathbb{E} \left[ \sum_{t=1}^T \mu \left( X_t(\alpha_{i_t} | \mathcal{F}_{t-1})^T \theta \right) \right] \leq 2\sqrt{(e-1)nT \log n},$$

where  $J = \{\alpha_1, \dots, \alpha_n\}$  is the tuning set of the hyper-parameter and  $|J| = n$ .

To bound Quantity (A), we note that we are not able to use any existing regret bound in the literature directly since the past information  $\mathcal{F}_{t-1}$  here is based on the sequence of arms pulled by our auto-tuning algorithm instead of the arms generated by using  $\alpha^*$  at each round, and the history would affect the update of bandit algorithms. We overcome this challenge by noticing that the consistency of  $\hat{\theta}_t$  plays a vital role in most of the proofs for (generalized) linear bandits, and this consistency could hold after a warm-up period or with large exploration rate. Therefore, we can expect a tight bound of the cumulative regret by using the same exploration parameter even under another line of observations  $\mathcal{F}_{t-1}$  with sufficient exploration. Another crux of proof is that the regret is usually related to  $\|x_t\|_{V_t^{-1}}$ , which can be similarly bounded after sufficient exploration. After we bound Quantity (A), combining Lemma 1, we get the following theorem.

**Theorem 1.** *Assume given the past information  $\mathcal{F}_{t-1}$  generated from our proposed algorithm for arm selection and the hyper-parameter to be used by the contextual bandits, the arm to be pulled follows a fixed distribution. For UCB and TS based generalized linear bandit algorithms with exploration hyper-parameters (LinUCB, UCB-GLM, LinTS, ect.), the regret bound of Algorithm 1 satisfies*

$$(1) \quad \mathbb{E}[R(T)] = \tilde{O}(T^{2/3}) + O(\sqrt{n(T - T_1) \log n}) \text{ given the warm-up length } T_1 = \tilde{O}(T^{2/3}).$$

(2) For UCB-based bandits, if the theoretical exploration parameter  $\alpha(T)$  is no larger than any element in  $J$ , then it holds that  $\mathbb{E}[R(T)] = \tilde{O}(\sqrt{T}) + O(\sqrt{nT \log n})$  with  $T_1 = 0$ .

(3) If  $\mathcal{A}_t$  is a convex set, and the smallest principal curvature in any neighborhood of the optimal vector  $x_{t,*} \in \mathcal{A}_t$  on  $\mathcal{A}_t$  can be lower bounded by some positive constant  $c$ , then  $\mathbb{E}[R(T)] = \tilde{O}(T^{4/7}) + O(\sqrt{n(T - T_1) \log n})$  after a warming-up period of length  $T_1 = O(T^{4/7})$ .

**Remark 1.** We could expect a similar result for TS-based bandit algorithms as in Theorem 1(2), and we offer an intuitive explanation in Appendix 4. Moreover, the conditions in Proposition 1(3) could be easily verified in many cases. For example, it holds when  $\mathcal{A}_t = \{x \in \mathbb{R}^d : \|x\| \leq a\}, \forall a > 0$ .

## 5 The Syndicated Bandits framework for selecting multiple hyper-parameters

There can be multiple hyper-parameters in the contextual bandit algorithm. For example, in linear bandit algorithms such as LinUCB [1, 17] and LinTS [5], exploration parameter  $\alpha$  and the regularization parameter of the ridge regression  $\lambda$  are both hyper-parameters to be tuned. In more recent contextual bandit works, there could be even more than two hyper-parameters. For example, NeuralUCB algorithm [28], which is proposed for the contextual bandit problems with a deep neural network model, has many tuning parameters such as the network width, network depth, step size for gradient descent, number of steps for gradient descent, as well as exploration parameter and regularization parameter  $\lambda$ , etc. Another example can be found in [14], where an efficient SGD-TS algorithm is proposed for generalized linear bandits, the number of tuning parameters is also more than two.

A naive strategy to auto-tune multiple hyper-parameters is to use Algorithm 1 and let the tuning set  $J$  contain all the possible combinations of the hyper-parameters. Assume there are in total  $L$  hyper-parameters  $\alpha^{(1)}, \alpha^{(2)}, \dots, \alpha^{(L)}$ . For all  $l \in [L]$ , if the tuning set for  $\alpha^{(l)}$  is defined as  $J_l = \{\alpha_1^{(l)}, \dots, \alpha_{n_l}^{(l)}\}$ , where  $n_l$  is the size of the corresponding tuning set. Then there are in total  $\prod_{l=1}^L n_l$  possible combinations. Based on Lemma 1 the

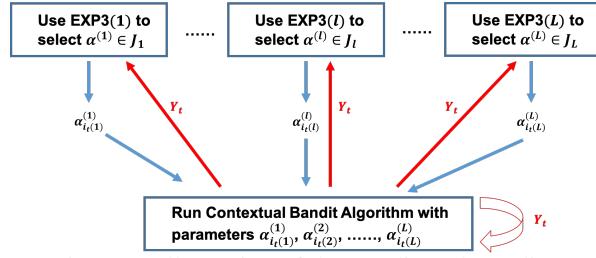


Figure 1: Illustration of the Syndicated Bandits.

extra regret paid to tune the hyper-parameters (Quantity (B)) is upper bounded by  $\tilde{O}(\sqrt{\prod_{l=1}^L n_l T})$ . Therefore, the aforementioned naive approach makes the regret increase exponentially with the number of tuning parameters. To mitigate this issue, we propose the Syndicated Bandits framework that can deal with multiple hyper-parameters while avoiding the exponential dependency on the number of tuning parameters in regret bounds.

We create  $L + 1$  bandit instances in this framework. In the bottom layer, the contextual bandit algorithm is used to decide which arm to pull. On top of the contextual bandit layer, there are  $L$  EXP3 bandits, denoted as  $\text{EXP3}(l)$  for  $l \in [L]$ . Each EXP3 algorithm is responsible for tuning one hyper-parameter only. At round  $t$ , if  $i_t(l)$  is the index of the hyper-parameters in  $J_l$  selected by the  $\text{EXP3}(l)$  bandit and the selected hyper-parameter is denoted as  $\alpha_{i_t(l)}^{(l)}$  for  $l \in [L]$ , then the contextual bandit algorithm in the bottom layer will use these parameters to make a decision and receive a reward based on the pulled arm. The reward is fed to all the  $L + 1$  bandits to update the components. Illustration of the algorithm and more details are presented in Figure 1 and Algorithm 2 in Appendix.

### 5.1 Regret analysis

At round  $t$ , given all the past information  $\mathcal{F}_{t-1}$ , denote  $a_t(\alpha_{j_1}^{(1)}, \dots, \alpha_{j_L}^{(L)} | \mathcal{F}_{t-1})$  as the arm pulled by the contextual bandit algorithm if the parameters are chosen as  $\alpha^{(l)} = \alpha_{j_l}^{(l)}$  for all  $l \in [L]$  and let  $X_t(\alpha_{j_1}^{(1)}, \dots, \alpha_{j_L}^{(L)} | \mathcal{F}_{t-1})$  be the corresponding feature vector. Recall that  $\mu(\cdot)$  is the reward-feature

model function, then for an arbitrary combination of hyper-parameters  $(\alpha_*^{(1)}, \dots, \alpha_*^{(L)})$ ,

$$\begin{aligned} \mathbb{E}[R(T)] &= \sum_{t=1}^{T_1} \mathbb{E} [\mu(x_{t,*}^T \theta) - \mu(X_t^T \theta)] + \sum_{t=T_1+1}^T \mathbb{E} [\mu(x_{t,*}^T \theta) - \mu(X_t(\alpha_*^{(1)}, \dots, \alpha_*^{(L)} | \mathcal{F}_{t-1})^T \theta)] \\ &+ \sum_{t=T_1}^T \mathbb{E} [\mu(X_t(\alpha_*^{(1)}, \dots, \alpha_*^{(L)} | \mathcal{F}_{t-1})^T \theta) - \mu(X_t(\alpha_{i_t(1)}^{(1)}, \alpha_*^{(2)}, \dots, \alpha_*^{(L)} | \mathcal{F}_{t-1})^T \theta)] \\ &+ \sum_{t=T_1}^T \mathbb{E} [\mu(X_t(\alpha_{i_t(1)}^{(1)}, \alpha_*^{(2)}, \dots, \alpha_*^{(L)} | \mathcal{F}_{t-1})^T \theta) - \mu(X_t(\alpha_{i_t(1)}^{(1)}, \alpha_{i_t(2)}^{(2)}, \alpha_*^{(3)}, \dots | \mathcal{F}_{t-1})^T \theta)] \\ &+ \dots + \sum_{t=T_1}^T \mathbb{E} [\mu(X_t(\alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(L-1)}^{(L-1)}, \alpha_*^{(L)} | \mathcal{F}_{t-1})^T \theta) - \mu(X_t(\alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(L)}^{(L)} | \mathcal{F}_{t-1})^T \theta)]. \end{aligned}$$

The first quantity represents the regret from pure exploration. The second quantity in the above decomposition is the regret of the contextual bandit algorithm that runs with the same hyper-parameters  $\alpha_*^{(1)}, \dots, \alpha_*^{(L)}$  under the past history  $\mathcal{F}_{t-1}$  generated from our tuning strategy every round. The next  $L$  quantities in the decomposition are the regret from tuning parameters in the EXP3 layer, which can be bounded using similar techniques in Lemma 1. However, the correlations between parameters are more complicated in the analysis now. Formally, we provide the following Theorem to guarantee the performance of the Syndicated Bandits framework. Proofs are deferred to the Appendix.

**Theorem 2.** *Assume given the past information  $\mathcal{F}_{t-1}$  and the hyper-parameters to be used by the contextual bandit algorithm at round  $t$ , the arm to be pulled by the contextual bandit algorithm follows a fixed distribution. Then the auto tuning method in Algorithm 2 with warm-up length  $T_1 = O(T^{2/3})$  has the following regret bound in general:*

$$\mathbb{E}[R(T)] \leq \tilde{O}(T^{2/3}) + O\left(\sum_{l=1}^L \sqrt{n_l(T - T_1) \log n_l}\right).$$

**Remark 2.** *Note this result avoids the exponential dependency on the number of hyper-parameters to be tuned in regret. When the hyper-parameters to be tuned are the exploration parameter  $\alpha$  and the regularization parameter  $\lambda$  of the (generalized) linear model, we also have the same conclusions as in Theorem 1(3). Please refer to Appendix A.3 for a formal statement and its proof.*

**Remark 3.** *Without any assumptions, Algorithm 2 has its regret dependent on  $d$  as  $O(d^3 + dT^{2/3})$  for both UCB and TS. In practice, usually  $d \ll T$ .*

## 6 Experimental results

We show by experiments that our proposed methods outperform various contextual bandit algorithm using the theoretical exploration parameter, as well as existing tuning methods. We compare different hyper-parameter selection methods in three popular contextual bandit algorithms, LinUCB [1, 17], LinTS [5] and UCB-GLM [18] with a logistic model. In practice, we set the warm-up length as  $T_1 = 0$  and tune both exploration parameters and regularization parameters. We compare the following hyper-parameter selection methods. **Theoretical-Explore** [1]: At round  $t$ , this method uses the true theoretical exploration parameter  $\alpha(t)$  defined in Equation 3; **OP** [9]: We make simple modifications of OPLINUCB to make it applicable to tune exploration parameters for LinUCB, LinTS and UCB-GLM; **Corral** [3]: This method uses the corraling idea to tune the exploration parameter only. **Corral-Combined** [3]: This method treats bandits with different combinations of the exploration parameter and regularization parameter  $\lambda$  as base model and uses the corraling idea to tune the configurations; **TL (Our work, Algorithm 1)**: This is our proposed Algorithm 1, where we use the two-layer bandit structure to tune the exploration parameter only; **TL-Combined (Our work, Algorithm 1)**: This method tunes both the exploration parameter  $\alpha$  and the regularization parameter  $\lambda$  using Algorithm 1 but with the tuning set containing all the possible combinations of  $\alpha$  and  $\lambda$ ; **Syndicated (Our work, Algorithm 2)**: This method keeps two separate tuning sets for  $\alpha$  and  $\lambda$  respectively. It uses the Syndicated Bandits in Algorithm 2.

We set the tuning set for exploration parameter  $\alpha$  as  $\{0, 0.01, 0.1, 1, 10\}$  and set the tuning set for regularization parameter  $\lambda$  as  $\{0.01, 0.1, 1\}$  in TL-Combined, Corral-Combined and Syndicated.

---

**Algorithm 2** The Syndicated Bandits Framework for Auto Tuning Multiple Hyper-parameters

---

**Input:** time horizon  $T$ , warm up length  $T_1$ , candidate hyper-parameter set  $\{J_l\}_{l=1}^L$ .

- 1: Randomly choose  $a_t \in [K]$  and record  $X_t, Y_t$  for  $t \in [T_1]$ .
- 2: Initialize exponential weights  $w_j^{(l)}(T_1 + 1) = 1$  for  $t = 1, j = 1, \dots, n_l$  and  $l = 1, \dots, L$ .
- 3: Initialize the parameters for all EXP3 layers as  $\beta_l = \min\left\{1, \sqrt{\frac{n_l \log n_l}{(e-1)T}}\right\}$ .
- 4: **for**  $t = (T_1 + 1)$  **to**  $T$  **do**
- 5:     Update probability distribution for pulling candidates in  $J_l$

$$p_j^{(l)}(t) = \frac{\beta_l}{n_l} + (1 - \beta_l) \frac{w_j^{(l)}(t)}{\sum_{i=1}^{n_l} w_i^{(l)}(t)}$$

- 6:      $i_t(l) \leftarrow j \in [n_l]$  with probability  $p_j^{(l)}(t)$  for all  $l = 1, \dots, L$ .
- 7:     Run the contextual bandit algorithm with hyper-parameters  $\alpha^{(l)} = \alpha_{i_t(l)}^{(l)}$  to pull an arm.
- 8:     Observe reward  $Y_t$  and update the components in contextual bandit algorithms.
- 9:     Update all  $L$  EXP3 bandits:  $\hat{y}_t^{(l)}(j) \leftarrow 0$  if  $j \neq i_t(l)$ . Otherwise,  $\hat{y}_t^{(l)}(j) \leftarrow Y_t / p_j^{(l)}(t)$ .
- 10:    For all  $l = 1, \dots, L$ , let  $w_j^{(l)}(t+1) = w_j^{(l)}(t) \times \exp\left(\frac{\beta_l}{n_l} \hat{y}_t^{(l)}(j)\right)$ .
- 11: **end for**

---

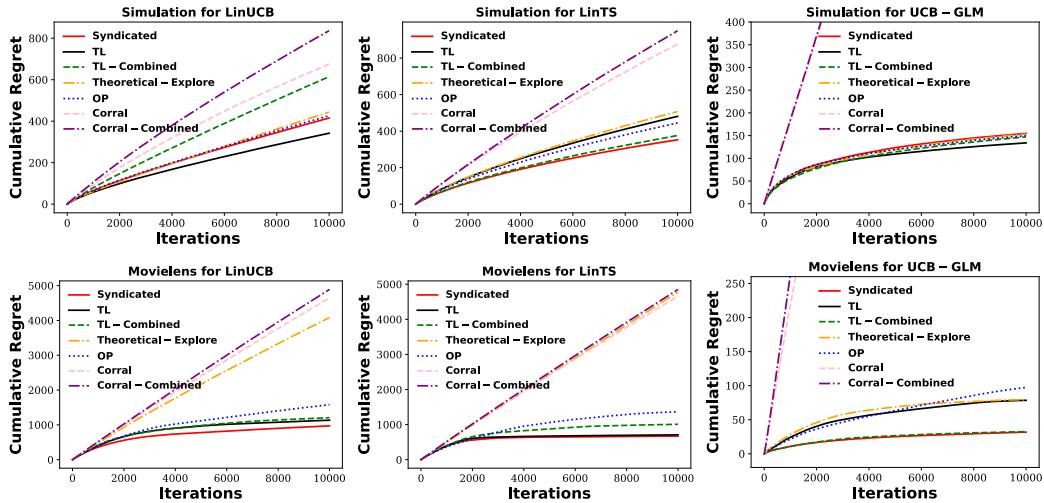


Figure 2: Comparison of hyper-parameters selection methods in LinUCB, LinTS and UCB-GLM.

For Theoretical-Explore, OP and TL, since they only tune the exploration parameter, we set the regularization parameter as  $\lambda = 1$ . In all the experiments below, the total number of rounds is  $T = 10,000$ . We run the comparisons on both simulations and the benchmark MovieLens 100K real datasets. Due to limited space, the descriptions of the dataset settings are deferred to Appendix A.4. Averaged results over 10 independently repeated experiments are reported below.

From Figure 2, we observe: 1) When tuning only one hyper-parameter (exploration parameter in our experiments), the proposed method outperforms previous tuning methods. Further, the theoretical exploration parameter does not perform well and it tends to be too conservative in practice, which is consistent with the results we show in Table I. 2) When tuning multiple hyper-parameters, previous methods do not apply. We found using the Syndicated Bandits framework usually outperforms TL-Combined and is significantly better than Corral-Combined method which has exponential regret with respect to number of tuning parameters. 3) Using Syndicated Bandits to tune multiple hyper-parameters usually outperforms tuning one parameter only. This demonstrates a practical need of auto-tuning multiple hyper-parameters in bandit algorithms. See Appendix for additional experiments on the tuning 3 hyper-parameters in SGD-TS [44].

## 7 Conclusion

In this paper, we propose a two-layer bandit structure for auto tuning the exploration parameter in contextual bandit algorithms, where the offline tuning is impossible. To further accommodate multiple hyper-parameters tuning tasks in contextual bandit algorithms with complicated models, we generalize our method to the Syndicated Bandits framework. This is the first framework that can auto-tune multiple hyper-parameters dynamically from observations in contextual bandit environment with theoretical regrets that avoids exponential dependency on the total number of hyper-parameters to be tuned. We show that our proposed algorithm can obtain  $\tilde{O}(T^{2/3})$  regret in general and has optimal  $\tilde{O}(\sqrt{T})$  regret for UCB-based algorithms when the all candidates in the tuning set is greater than the theoretical exploration parameter. Our work is general enough to handle the tuning tasks in many contextual bandit algorithms. Experimental results also validate the effectiveness of our proposed work.

## Acknowledgments and Disclosure of Funding

We are grateful for the insightful comments from the anonymous reviewers and area chair. This work was partially supported by the National Science Foundation under grants CCF-1934568, DMS-1811405, DMS-1811661, DMS-1916125, DMS-2113605, DMS-2210388, IIS-2008173 and IIS-2048280. CJH is also supported by Samsung, Google, Sony and the Okawa Foundation.

## References

- [1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- [2] Marc Abeille and Alessandro Lazaric. Linear thompson sampling revisited. In *Artificial Intelligence and Statistics*, pages 176–184. PMLR, 2017.
- [3] Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E Schapire. Corralling a band of bandit algorithms. In *Conference on Learning Theory*, pages 12–38. PMLR, 2017.
- [4] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on learning theory*, pages 39–1, 2012.
- [5] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- [6] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [7] 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.
- [8] Ilija Bogunovic, Arpan Losalka, Andreas Krause, and Jonathan Scarlett. Stochastic linear bandits robust to adversarial attacks. In *International Conference on Artificial Intelligence and Statistics*, pages 991–999. PMLR, 2021.
- [9] Djallel Bouneffouf and Emmanuelle Claeys. Hyper-parameter tuning for the contextual bandit. *arXiv preprint arXiv:2005.02209*, 2020.
- [10] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257, 2011.
- [11] Wang Chi Cheung, David Simchi-Levi, and Ruihao Zhu. Learning to optimize under non-stationarity. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1079–1087, 2019.
- [12] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pages 844–853. PMLR, 2017.

[13] Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *COLT*, pages 355–366, 2008.

[14] Qin Ding, Cho-Jui Hsieh, and James Sharpnack. An efficient algorithm for generalized linear bandit: Online stochastic gradient descent and thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 1585–1593. PMLR, 2021.

[15] Sarah Filippi, Olivier Cappe, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *NIPS*, volume 23, pages 586–594, 2010.

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

[17] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.

[18] Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning*, pages 2071–2080. PMLR, 2017.

[19] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research*, 18(1):6765–6816, 2017.

[20] Eric M Schwartz, Eric T Bradlow, and Peter S Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017.

[21] Xuedong Shang, Emilie Kaufmann, and Michal Valko. A simple dynamic bandit algorithm for hyper-parameter tuning. In *Workshop on Automated Machine Learning at International Conference on Machine Learning*, Long Beach, United States, June 2019. AutoML@ICML 2019 - 6th ICML Workshop on Automated Machine Learning.

[22] Amr Sharaf and Hal Daumé III. Meta-learning for contextual bandit exploration. *arXiv preprint arXiv:1901.08159*, 2019.

[23] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

[24] Jingkang Wang, Mengye Ren, Ilija Bogunovic, Yuwen Xiong, and Raquel Urtasun. Cost-efficient online hyperparameter optimization. *arXiv preprint arXiv:2101.06590*, 2021.

[25] Michael Woodroffe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.

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

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

[28] Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 2020.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
  - (b) Did you describe the limitations of your work? **[N/A]**

- (c) Did you discuss any potential negative societal impacts of your work? [N/A]
- (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]

2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix.
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We include the code in the submitted supplementary materials.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 6
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Table 1
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes]
  - (b) Did you mention the license of the assets? [Yes] See Appendix for the dataset explanation
  - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include the code in the submitted supplementary materials.
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A Appendix

### A.1 Proof of Lemma 1

*Proof.* The proof techniques basically follows [7]. However, since the EXP3 layer and contextual bandit layer are coupled, the result in [7] cannot be directly applied to show our result. We make modifications of the proofs in [7] below.

We first reload some notations in this proof: at time  $t$  we are given all the previous information  $\mathcal{F}_{t-1}$  generated from using our auto-tune framework shown in Algorithm 1, and then pull an arm according to some exploration hyper-parameter  $\alpha$ . Therefore, for convenience we could safely omit  $\mathcal{F}_{t-1}$  here, and denote  $a_t(\alpha) := a_t(\alpha|\mathcal{F}_{t-1})$  and  $X_t(\alpha) := X_t(\alpha|\mathcal{F}_{t-1})$  as the arm pulled and its corresponding feature vector at round  $t$ . Furthermore, if arm  $a_t(\alpha_j)$  is pulled at round  $t$ , we define the corresponding mean reward as  $\mu_t(\alpha_j) = \mu(X_t(\alpha_j)^T \theta)$ . The corresponding observed sample reward is  $y_t(\alpha_j) = \mu_t(\alpha_j) + \epsilon_{t,j}$ , where  $\epsilon_{t,j}$  denotes the hypothetical random noise at round  $t$  if arm  $a_t(\alpha_j)$  is pulled. Note that  $\epsilon_t = \epsilon_{t,i_t}$  since  $a_t(\alpha_{i_t})$  is the arm pulled by our algorithm and  $\epsilon_t$  is the associated random noise. By definition,  $Y_t = y_t(\alpha_{i_t})$ . From the definition of  $\hat{y}_t(j)$  in Algorithm 1, we have  $\hat{y}_t(j) = y_t(\alpha_j)/p_j(t)$  if  $j = i_t$ . Otherwise  $\hat{y}_t(j) = 0$ . Then  $w_j(t+1) = w_j(t) \exp(\frac{\beta}{n} \hat{y}_t(i))$  according to Algorithm 1.

Given all the information in the past  $\mathcal{F}_{t-1}$ ,  $(\hat{\theta}_t, V_t, p_j(t), w_j(t))$  are fixed. Since  $0 \leq y_t(\alpha_j) \leq 1$ , we have

$$\mathbb{E} \left[ \sum_{i=1}^n p_i(t) \hat{y}_t(i) \mid \mathcal{F}_{t-1} \right] = \mathbb{E} \left[ p_{i_t}(t) \frac{y_t(\alpha_{i_t})}{p_{i_t}(t)} \mid \mathcal{F}_{t-1} \right] = \mathbb{E} [\mu_t(\alpha_{i_t}) \mid \mathcal{F}_{t-1}] \quad (4)$$

$$\begin{aligned} \mathbb{E} \left[ \sum_{i=1}^n p_i(t) \hat{y}_t(i)^2 \mid \mathcal{F}_{t-1} \right] &= \mathbb{E} \left[ p_{i_t}(t) \frac{y_t(\alpha_{i_t})}{p_{i_t}(t)} \hat{y}_t(i_t) \mid \mathcal{F}_{t-1} \right] = \mathbb{E} [y_t(\alpha_{i_t}) \hat{y}_t(i_t) \mid \mathcal{F}_{t-1}] \\ &\leq \mathbb{E} [\hat{y}_t(i_t) \mid \mathcal{F}_{t-1}] = \mathbb{E} \left[ \sum_{i=1}^n \hat{y}_t(i) \mid \mathcal{F}_{t-1} \right] \end{aligned} \quad (5)$$

$$= \sum_{i=1}^n \mathbb{E} [\mathbb{E} [\hat{y}_t(i) \mid \sigma(\mathcal{F}_{t-1}, \epsilon_{t,i}, a_t(\alpha_i))] \mid \mathcal{F}_{t-1}] \quad (6)$$

$$= \sum_{i=1}^n \mathbb{E} [y_t(\alpha_i) \mid \mathcal{F}_{t-1}] \quad (7)$$

$$= \sum_{i=1}^n \mathbb{E} [\mu_t(\alpha_i) \mid \mathcal{F}_{t-1}]. \quad (8)$$

Equation 5 holds since  $\hat{y}_t(i) \neq 0$  only when  $i = i_t$ . In Equation 6,  $\sigma(\mathcal{F}_{t-1}, \epsilon_{t,i}, a_t(\alpha_i))$  is the smallest  $\sigma$ -algebra induced by  $\mathcal{F}_{t-1}$ ,  $\epsilon_{t,i}$ , and  $a_t(\alpha_i)$ . Equation 7 holds since  $\hat{y}_t(i) = y_t(\alpha_i)/p_i(t) \mathbb{1}(i = i_t)$ . Meanwhile, since given the hyper-parameter to be used at round  $t$  as  $\alpha_i$ , the arm to be pulled  $a_t(\alpha_i)$  follows a fixed distribution and does not affect the distribution of  $i_t$ , so  $i = i_t$  is still with probability  $p_i(t)$ . Now we are ready to use the above results to prove the lemma. Define  $W_t = \sum_{i=1}^n w_i(t)$ . We find the lower bound and upper bound of  $\mathbb{E}[\log \frac{W_{t+1}}{W_t}]$  below.

**Lower bound.** Since  $w_i(1) = 1$  for all  $i$ ,  $\mathbb{E}[\log \frac{W_{t+1}}{W_t}] \geq \mathbb{E}[\log w_i(T+1)] - \log n$  for all  $i \in [n]$ .

We take a look at  $\mathbb{E} \left[ \log \frac{w_i(t+1)}{w_i(t)} \right]$  below.

$$\begin{aligned} \mathbb{E} \left[ \log \frac{w_i(t+1)}{w_i(t)} \mid \mathcal{F}_{t-1} \right] &= \mathbb{E} \left[ \log \left[ \frac{w_i(t)}{w_i(t)} \exp \left( \frac{\beta}{n} \hat{y}_t(i) \right) \right] \mid \mathcal{F}_{t-1} \right] \\ &= \mathbb{E} \left[ \frac{\beta}{n} \hat{y}_t(i) \mid \mathcal{F}_{t-1} \right] = \mathbb{E} \left[ \frac{\beta}{n} y_t(i) \mid \mathcal{F}_{t-1} \right] = \mathbb{E} \left[ \frac{\beta}{n} \mu_t(\alpha_i) \mid \mathcal{F}_{t-1} \right] \end{aligned}$$

The third “=” in the above is due to the same reason as in Equation 7. Take expectation on both sides and sum over  $t$ , we get

$$\mathbb{E}[\log w_i(T+1)] = \frac{\beta}{n} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_i)]$$

Therefore, for all  $i = 1, \dots, n$ ,

$$\mathbb{E} \left[ \log \frac{W_{T+1}}{W_1} \right] \geq \frac{\beta}{n} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_i)] - \log n. \quad (9)$$

**Upper bound.** On the other hand, let's look at  $\mathbb{E}[\log \frac{W_{t+1}}{W_t}]$ :

$$\begin{aligned} \mathbb{E} \left[ \log \frac{W_{t+1}}{W_t} \mid \mathcal{F}_{t-1} \right] &= \mathbb{E} \left[ \log \sum_{i=1}^n \frac{w_i(t+1)}{W_t} \mid \mathcal{F}_{t-1} \right] = \mathbb{E} \left[ \log \sum_{i=1}^n \frac{w_i(t)}{W_t} \exp \left( \frac{\beta}{n} \hat{y}_t(i) \right) \mid \mathcal{F}_{t-1} \right] \\ &= \mathbb{E} \left[ \log \sum_{i=1}^n \frac{p_i(t) - \frac{\beta}{n}}{1 - \beta} \exp \left( \frac{\beta}{n} \hat{y}_t(i) \right) \mid \mathcal{F}_{t-1} \right] \quad \text{definition of } p_i(t) \\ &\leq \mathbb{E} \left[ \log \sum_{i=1}^n \frac{p_i(t) - \frac{\beta}{n}}{1 - \beta} \left( 1 + \frac{\beta}{n} \hat{y}_t(i) + \frac{(e-2)\beta^2}{n^2} \hat{y}_t(i)^2 \right) \mid \mathcal{F}_{t-1} \right] \\ &\leq \mathbb{E} \left[ \log \left( 1 + \sum_{i=1}^n \left[ \frac{\beta}{n(1-\beta)} p_i(t) \hat{y}_t(i) + \frac{(e-2)\beta^2}{n^2(1-\beta)} p_i(t) \hat{y}_t(i)^2 \right] \right) \mid \mathcal{F}_{t-1} \right] \\ &\leq \mathbb{E} \left[ \sum_{i=1}^n \left( \frac{\beta}{n(1-\beta)} p_i(t) \hat{y}_t(i) + \frac{(e-2)\beta^2}{n^2(1-\beta)} p_i(t) \hat{y}_t(i)^2 \mid \mathcal{F}_{t-1} \right) \right] \\ &\leq \frac{\beta}{n(1-\beta)} \mathbb{E}[\mu_t(\alpha_{i_t}) \mid \mathcal{F}_{t-1}] + \frac{(e-2)\beta^2}{n^2(1-\beta)} \sum_{i=1}^n \mathbb{E}[\mu_t(\alpha_i) \mid \mathcal{F}_{t-1}]. \end{aligned}$$

The first inequality in the above holds since  $e^x \leq 1 + x + (e-2)x^2$  for  $x \in [0, 1]$ . Here, we have  $0 \leq \frac{\beta}{n} \hat{y}_t(i) \leq 1$  because  $p_i(t) \geq \frac{\beta}{n}$  and  $0 \leq y_t(\alpha_i) \leq 1$ . The third inequality “ $\leq$ ” in the above holds since  $\log(1+x) \leq x$  when  $x \geq 0$ . The last inequality is from Equation 4, 8. Take another expectation on both sides, we get

$$\mathbb{E} \left[ \log \frac{W_{t+1}}{W_t} \right] \leq \frac{\beta}{n(1-\beta)} \mathbb{E}[\mu_t(\alpha_{i_t})] + \frac{(e-2)\beta^2}{n^2(1-\beta)} \sum_{i=1}^n \mathbb{E}[\mu_t(\alpha_i)]$$

By summing the above over  $t$ , we have

$$\mathbb{E} \left[ \log \frac{W_{T+1}}{W_1} \right] \leq \frac{\beta}{n(1-\beta)} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_{i_t})] + \frac{(e-2)\beta^2}{n^2(1-\beta)} \sum_{t=1}^T \sum_{i=1}^n \mathbb{E}[\mu_t(\alpha_i)] \quad (10)$$

Combining the lower bound (Equation 9) and upper bound (Equation 10) of  $\mathbb{E} \left[ \log \frac{W_{T+1}}{W_1} \right]$ , we get for every  $i = 1, \dots, n$ ,

$$\frac{\beta}{n} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_i)] - \log n \leq \frac{\beta}{n(1-\beta)} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_{i_t})] + \frac{(e-2)\beta^2}{n^2(1-\beta)} \sum_{t=1}^T \sum_{i=1}^n \mathbb{E}[\mu_t(\alpha_i)] \quad (11)$$

Let

$$G_{\max} = \max_{i \in [n]} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_i)]$$

Since Equation 11 holds for any  $i$ , we have

$$\frac{\beta}{n} G_{\max} - \log n \leq \frac{\beta}{n(1-\beta)} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_{i_t})] + \frac{(e-2)\beta^2}{n(1-\beta)} G_{\max} \quad (12)$$

Equation 12 can be further simplified as

$$G_{\max} - \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_{i_t})] \leq (e-1)\beta G_{\max} + \frac{(1-\beta)n \log n}{\beta}$$

Since we choose  $\beta = \min\left\{1, \sqrt{\frac{n \log n}{(e-1)T}}\right\}$  and note that  $G_{\max} \leq T$ , we get

$$G_{\max} - \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_{i_t})] \leq 2\sqrt{(e-1)Tn \log n} = \tilde{O}(\sqrt{nT}).$$

□

## A.2 Proof of Theorem 1

To bound the cumulative regret, we only need to bound Quantity (A) and then combine the results in Lemma 1. In the following, we first list some useful lemmas for bounding Quantity (A) for completeness.

### A.2.1 Useful Lemmas

**Lemma 2** (Proposition 1 in [18]). *Define  $V_{n+1} = \sum_{t=1}^n X_t X_t^T$ , where  $X_t$  is drawn IID from some distribution in unit ball  $\mathbb{B}^d$ . Furthermore, let  $\Sigma := E[X_t X_t^T]$  be the second moment matrix, let  $B, \delta_2 > 0$  be two positive constants. Then there exists positive, universal constants  $C_1$  and  $C_2$  such that  $\lambda_{\min}(V_{n+1}) \geq B$  with probability at least  $1 - \delta_2$ , as long as*

$$n \geq \left( \frac{C_1 \sqrt{d} + C_2 \sqrt{\log(1/\delta_2)}}{\lambda_{\min}(\Sigma)} \right)^2 + \frac{2B}{\lambda_{\min}(\Sigma)}.$$

**Lemma 3** (Theorem 2 in [1]). *For any  $\delta < 1$ , under our problem setting in Section 3, it holds that for all  $t > 0$ ,*

$$\begin{aligned} \|\hat{\theta}_t - \theta^*\|_{V_t} &\leq \beta_t(\delta), \\ \forall x \in \mathbb{R}^d, |x^\top (\hat{\theta}_t - \theta^*)| &\leq \|x\|_{V_t^{-1}} \beta_t(\delta), \end{aligned}$$

with probability at least  $1 - \delta$ , where

$$\beta_t(\delta) = \sigma \sqrt{\log \left( \frac{(\lambda + t)^d}{\delta^2 \lambda^d} \right)} + \sqrt{\lambda} S.$$

In this subsection we denote  $\alpha^*(\delta) := \beta_T(\delta)$ .

**Lemma 4** ([15]). *Let  $\lambda > 0$ , and  $\{x_i\}_{i=1}^t$  be a sequence in  $\mathbb{R}^d$  with  $\|x_i\| \leq 1$ , then we have*

$$\begin{aligned} \sum_{s=1}^t \|x_s\|_{V_s^{-1}}^2 &\leq 2 \log \left( \frac{\det(V_{t+1})}{\det(\lambda I)} \right) \leq 2d \log \left( 1 + \frac{t}{\lambda} \right), \\ \sum_{s=1}^t \|x_s\|_{V_s^{-1}} &\leq \sqrt{t \left( \sum_{s=1}^t \|x_s\|_{V_s^{-1}}^2 \right)} \leq \sqrt{2dt \log \left( 1 + \frac{t}{\lambda} \right)}. \end{aligned}$$

**Lemma 5** ([5]). *For a Gaussian random variable  $Z$  with mean  $m$  and variance  $\sigma^2$ , for any  $z \geq 1$ ,*

$$P(|Z - m| \geq z\sigma) \leq \frac{1}{\sqrt{\pi}z} e^{-z^2/2}.$$

### A.2.2 Formal Proof

*Proof.* Here we would use LinUCB and LinTS for the detailed proof, and note that regret bound of all other UCB and TS algorithms could be similarly deduced. First we will show the proof of Theorem I for the UCB-typed algorithm (e.g. LinUCB):

(1). Since  $\alpha^*$  in our regret decomposition could be arbitrary element in  $J$ , here we simply take  $\alpha^* = \min_{\alpha \in J} \alpha$ . For LinUCB, since the Lemma 3 holds for any sequence  $(x_1, \dots, x_t)$ , and hence we have that with probability at least  $1 - \delta$ ,

$$\|\hat{\theta} - \theta\|_{V_t} \leq \beta_t(\delta) \leq \alpha(t, \delta),$$

where  $\alpha(T, \delta)$  is the theoretical optimal exploration rate at round  $t$  we denoted in Eqn. (3) with probability parameter  $\delta$ . And we would omit  $\delta$  for simplicity. Recall that for  $t > T_1$ , we denote the feature vector pulled at round  $t$  as  $X_t$ , i.e.

$$X_t = \underset{x \in \mathcal{A}_t}{\operatorname{argmax}} x^\top \hat{\theta}_t + \alpha_{i_t} \|x\|_{V_t^{-1}}, \quad X_t = X_t(\alpha_{i_t} | \mathcal{F}_{t-1}).$$

And we also define  $\tilde{X}_t = X_t(\alpha^* | \mathcal{F}_{t-1})$ , i.e.

$$\tilde{X}_t = \underset{x \in \mathcal{A}_t}{\operatorname{argmax}} x^\top \hat{\theta}_t + \alpha^* \|x\|_{V_t^{-1}}.$$

And it turns out that the Quantity (A) can be represented by

$$\mathbb{E} \left[ \sum_{t=T_1+1}^T (\mu(x_{t,*}^\top \theta) - \mu(X_t(\alpha^* | \mathcal{F}_{t-1})^\top \theta)) \right] = \mathbb{E} \left[ \sum_{t=T_1+1}^T (\mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta)) \right].$$

According to the proof of LinUCB we could similarly argue that

$$\begin{aligned} x_{t,*}^\top \theta - \tilde{X}_t^\top \theta &\leq \alpha^* \left( \|\tilde{X}_t\|_{V_t^{-1}} - \|\tilde{x}_{t,*}\|_{V_t^{-1}} \right) + \|x_{t,*} - \tilde{X}_t\|_{V_t^{-1}} \|\hat{\theta}_t - \theta\|_{V_t} \\ &\leq (\alpha^* + \alpha(T)) \|\tilde{X}_t\|_{V_t^{-1}} + \alpha(T) \|x_{t,*}\|_{V_t^{-1}}. \end{aligned}$$

In conclusion, we have that

$$\sum_{t=T_1+1}^T (\mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta)) = \tilde{O} \left( \sum_{t=T_1+1}^T \|\tilde{X}_t\|_{V_t^{-1}} + \sum_{t=T_1+1}^T \|x_{t,*}\|_{V_t^{-1}} \right).$$

By Lemma 2 and choosing  $T_1 = T^{2/3}$ , it holds that,

$$\sum_{t=T_1+1}^T \|x_{t,*}\|_{V_t^{-1}}, \quad \sum_{t=T_1+1}^T \|\tilde{X}_t\|_{V_t^{-1}} = O(T \times T^{-1/3}) = O(T^{2/3}).$$

Note we could use this procedure to bound the regret for other UCB bandit algorithms, since most of the proofs for generalized linear bandits are closely related to the rate of  $\sum_{t=T_1+1}^T \|\tilde{X}_t\|_{V_t^{-1}}$ . Finally, the cost of pure exploration is also of scale  $\tilde{O}(T^{2/3})$ , which concludes the proof.

(2). Here we simply take  $\alpha^* = \min_{\alpha \in J} \alpha$ . We also use LinUCB as an example here since other UCB-based algorithms with exploration hyper-parameters could be identically bounded. Based on the definition of  $X_t$  and  $\tilde{X}_t$ , we have that,

$$\begin{aligned} X_t^\top \hat{\theta}_t + \alpha_{i_t} \|X_t\|_{V_t^{-1}} &= X_t^\top \hat{\theta}_t + (\alpha_{i_t} - \alpha^*) \|X_t\|_{V_t^{-1}} + \alpha^* \|X_t\|_{V_t^{-1}} \\ &\geq \tilde{X}_t^\top \hat{\theta}_t + (\alpha_{i_t} - \alpha^*) \|\tilde{X}_t\|_{V_t^{-1}} + \alpha^* \|\tilde{X}_t\|_{V_t^{-1}} \\ &\geq X_t^\top \hat{\theta}_t + (\alpha_{i_t} - \alpha^*) \|\tilde{X}_t\|_{V_t^{-1}} + \alpha^* \|X_t\|_{V_t^{-1}}, \end{aligned}$$

which implies that

$$(\alpha_{i_t} - \alpha^*) \|X_t\|_{V_t^{-1}} \geq (\alpha_{i_t} - \alpha^*) \|\tilde{X}_t\|_{V_t^{-1}}.$$

Since we have that  $\alpha_{i_t} \geq \alpha^*$ , and when  $\alpha_{i_t} > \alpha^*$  it holds that

$$\|X_t\|_{V_t^{-1}} \geq \|\tilde{X}_t\|_{V_t^{-1}}, \quad \forall t > 0. \quad (13)$$

On the other hand, when  $\alpha_{i_t} = \alpha^*$  it holds that  $X_t = \tilde{X}_t$ , which consequently implies that

$$\|X_t\|_{V_t^{-1}} = \|\tilde{X}_t\|_{V_t^{-1}}, \quad \forall t > 0.$$

According to the proof of LinUCB we could similarly argue that

$$\begin{aligned} x_{t,*}^\top \theta - \tilde{X}_t^\top \theta &\leq \alpha^* \left( \|\tilde{X}_t\|_{V_t^{-1}} - \|\tilde{x}_{t,*}\|_{V_t^{-1}} \right) + \|x_{t,*} - \tilde{X}_t\|_{V_t^{-1}} \|\hat{\theta}_t - \theta\|_{V_t} \\ &\leq 2\alpha^* \|\tilde{X}_t\|_{V_t^{-1}}, \end{aligned}$$

since  $\alpha(T) \leq \alpha^*$ . Therefore, we have

$$\sum_{t=T_1+1}^T \left( \mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta) \right) \leq 2\alpha^* \sum_{t=T_1+1}^T \|\tilde{X}_t\|_{V_t^{-1}} \leq 2\alpha^* \sum_{t=T_1+1}^T \|X_t\|_{V_t^{-1}} = \tilde{O}(\sqrt{T}).$$

By taking  $T_1 = \tilde{O}(\sqrt{T})$ , we know the cumulative regret of the warm-up period is also of order  $\tilde{O}(\sqrt{T})$ , which concludes the proof.

**Remark 4.** (1) Intuitively, we can deduce Eqn. (13) by choosing  $\alpha^* = \min_{\alpha \in J} \alpha$ , i.e.  $\alpha^*$  is no larger than any exploration hyper-parameter candidate since the best feature vector solved in UCB algorithms tends to have larger value of  $\|\cdot\|_{V_t^{-1}}$  at time  $t$  if we enlarge  $\alpha$ . In other words, under larger  $\alpha$  we would more likely to choose arm with greater uncertainty quantified by the value of  $\|\cdot\|_{V_t^{-1}}$ . (2) On the other hand, for TS bandit algorithms we would expect the similar result: the feature vectors of superior arms tend to have smaller value of  $\|\cdot\|_{V_t^{-1}}$  since the value of  $\|\cdot\|_{V_t^{-1}}$  depicts the standard deviation of the feature vector. And the direction of the optimal arm should be frequently explored in the long run and hence its standard deviation is expected to be smaller than other inferior arms. By enlarging  $\alpha$ , we would have more chance to choose those sub-optimal arm with larger standard deviation and smaller estimated reward, which means results in Eqn. (13) could happen with high probability.

(3). Here we would use LinUCB for the detailed proof, and note that regret bound of all other UCB algorithms could be similarly deduced. W.l.o.g. we take  $\alpha^* = \min_{\alpha \in J} \alpha$

For LinUCB, since the Lemma 3 holds for any sequence  $(x_1, \dots, x_t)$ , and hence we have that with probability at least  $1 - \delta$ ,

$$\|\hat{\theta} - \theta\|_{V_t} \leq \beta_t(\delta) \leq \alpha(t, \delta).$$

And we would omit  $\delta$  for simplicity. Recall that for  $t > T_1$ , we denote the feature vector pulled at round  $t$  as  $X_t$ , i.e.

$$X_t = \underset{x \in \mathcal{A}_t}{\operatorname{argmax}} x^\top \hat{\theta}_t + \alpha_{i_t} \|x\|_{V_t^{-1}}, \quad X_t = X_t(\alpha_{i_t} | \mathcal{F}_{t-1}).$$

And we also define  $\tilde{X}_t = X_t(\alpha^* | \mathcal{F}_{t-1})$ , i.e.

$$\tilde{X}_t = \underset{x \in \mathcal{A}_t}{\operatorname{argmax}} x^\top \hat{\theta}_t + \alpha^* \|x\|_{V_t^{-1}}.$$

And it turns out that the Quantity (A) can be represented by

$$\mathbb{E} \left[ \sum_{t=T_1+1}^T \left( \mu(x_{t,*}^\top \theta) - \mu(X_t(\alpha^* | \mathcal{F}_{t-1})^\top \theta) \right) \right] = \mathbb{E} \left[ \sum_{t=T_1+1}^T \left( \mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta) \right) \right].$$

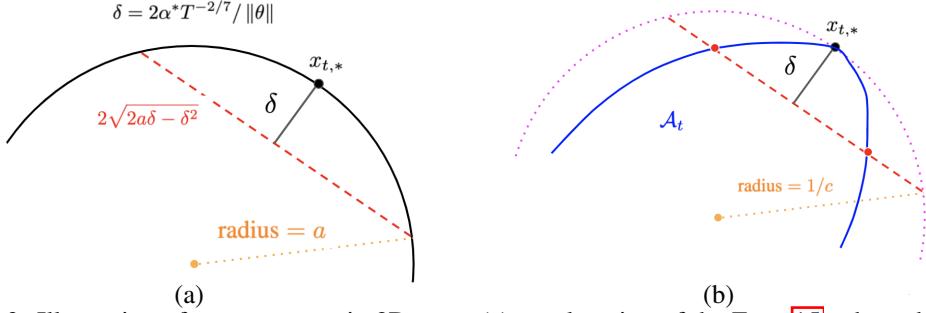


Figure 3: Illustration of our argument in 2D case: (a). explanation of the Eqn. 15, where the red line denotes the maximum distance between  $X_t$  and  $X_t^*$ ; (b). visualization on how to cover the neighborhood of  $x_{t,*}$  on  $\mathcal{A}_t$ , where the blue line denotes the boundary of  $\mathcal{A}_t$  and the pink dashed circle is the outer cover with radius  $1/c$ . In this case, the length of red line gives an upper bound of the maximum distance between  $X_t$  and  $X_t^*$ .

Note the selection of  $a_t$  in LinUCB implies that

$$x_{t,*}^\top \hat{\theta}_t + \alpha_{i_t} \|x_{t,*}\|_{V_t^{-1}} \leq X_t^\top \hat{\theta}_t + \alpha_{i_t} \|X_t\|_{V_t^{-1}}.$$

Therefore, we have

$$\begin{aligned} X_t^\top \hat{\theta}_t + \alpha_{i_t} \|X_t\|_{V_t^{-1}} &\geq x_{t,*}^\top \theta + \alpha_{i_t} \|x_{t,*}\|_{V_t^{-1}} + x_{t,*}^\top (\hat{\theta}_t - \theta) \\ &\geq x_{t,*}^\top \theta + \alpha_{i_t} \|x_{t,*}\|_{V_t^{-1}} - \|x_{t,*}\|_{V_t^{-1}} \|\hat{\theta}_t - \theta\|_{V_t} \\ &\geq x_{t,*}^\top \theta + (\alpha_{i_t} - \alpha(T)) \|x_{t,*}\|_{V_t^{-1}}. \end{aligned} \quad (14)$$

Therefore, it holds that,

$$\begin{aligned} X_t^\top \theta &\geq x_{t,*}^\top \theta + (\alpha_{i_t} - \alpha(T)) \|x_{t,*}\|_{V_t^{-1}} - \alpha_{i_t} \|X_t\|_{V_t^{-1}} + X_t^\top (\theta - \hat{\theta}_t), \\ (x_{t,*} - X_t)^\top \theta &\leq (\alpha(T) + \alpha_{i_t}) \|X_t\|_{V_t^{-1}} + (\alpha(T) - \alpha_{i_t}) \|x_{t,*}\|_{V_t^{-1}}, \end{aligned}$$

By Lemma 2, we have as long as  $T_1 = O(T^{4/7})$ , it holds that

$$(x_{t,*} - X_t)^\top \theta \leq 2\alpha(T)T^{-2/7}, \quad t > T_1.$$

Similarly, we could also deduce that

$$(x_{t,*} - \tilde{X}_t)^\top \theta \leq 2\alpha(T)T^{-2/7}, \quad t > T_1.$$

Firstly, we take  $\mathcal{A}_t = \{x : \|x\| \leq a^2\}$ ,  $a > 0$  for example, then it holds that  $x_{t,*} = \theta / \|\theta\|$ , and consequently

$$\|x_{t,*} - X_t\|, \|x_{t,*} - \tilde{X}_t\| \leq \sqrt{4a\alpha(T)T^{-2/7} / \|\theta\|} = O(\sqrt{\alpha(T)}T^{-1/7}).$$

Please refer to Figure 3 (a) for a 2D visual explanation, and similar argument could be made for higher dimension cases. And this implies that

$$\|X_t - \tilde{X}_t\| = O(\sqrt{\alpha(T)}T^{-1/7}). \quad (15)$$

Generally, if  $\mathcal{A}_t$  is some convex set, and we know there exists a small neighborhood of the optimal feature vector  $x_{t,*} \in \mathcal{A}_t$  such that the (sectional) principal curvature in this neighborhood can be lowered bounded by some positive constant  $c > 0$ . Then we can cover this neighborhood by a  $d$ -dimensional sphere with radius  $1/c$  (Figure 3 (b) for 2D visualization), and hence we could similarly deduce the above result. Note that for the example  $\mathcal{A}_t = \{x : \|x\| \leq a\}$ ,  $a > 0$ , all the principal curvatures are equal to  $1/a$  anywhere on this sphere, and hence it is a special case. The rest of argument is based on the proof outline of UCB bandits. According to the proof of LinUCB we could similarly argue that

$$\begin{aligned} x_{t,*}^\top \theta - \tilde{X}_t^\top \theta &\leq \alpha^* \left( \|\tilde{X}_t\|_{V_t^{-1}} - \|x_{t,*}\|_{V_t^{-1}} \right) + \|x_{t,*} - \tilde{X}_t\|_{V_t^{-1}} \|\hat{\theta}_t - \theta\|_{V_t} \\ &\leq (\alpha^* + \alpha(T)) \|\tilde{X}_t\|_{V_t^{-1}} + \alpha(T) \|x_{t,*}\|_{V_t^{-1}}. \end{aligned}$$

In conclusion, we have that

$$\sum_{t=T_1+1}^T \left( \mu(x_{t,*}^T \theta) - \mu(\tilde{X}_t^T \theta) \right) = \tilde{O} \left( \sum_{t=T_1+1}^T \left\| \tilde{X}_t \right\|_{V_t^{-1}} + \|x_{t,*}\|_{V_t^{-1}} \right). \quad (16)$$

Note that we have that,

$$\begin{aligned} \sum_{t=T_1+1}^T \left\| \tilde{X}_t \right\|_{V_t^{-1}} &\leq \sum_{t=T_1+1}^T \|X_t\|_{V_t^{-1}} + \sum_{t=T_1+1}^T \left\| X_t - \tilde{X}_t \right\|_{V_t^{-1}} \\ \sum_{t=T_1+1}^T \|x_{t,*}\|_{V_t^{-1}} &\leq \sum_{t=T_1+1}^T \|X_t\|_{V_t^{-1}} + \sum_{t=T_1+1}^T \|X_t - x_{t,*}\|_{V_t^{-1}}, \end{aligned}$$

where the first quantity could be easily bounded by Lemma 4, i.e.

$$\sum_{t=T_1+1}^T \|X_t\|_{V_t^{-1}} \leq \sqrt{2dT \log \left( 1 + \frac{T}{\lambda} \right)} = \tilde{O}(\sqrt{T}).$$

And the second quantity could be bounded as

$$\sum_{t=T_1+1}^T \left\| X_t - \tilde{X}_t \right\|_{V_t^{-1}} \leq \sum_{t=T_1+1}^T \left\| X_t - \tilde{X}_t \right\| \sqrt{\lambda_{\min}(V_t^{-1})} \lesssim \sum_{t=T_1+1}^T \sqrt{\alpha^*} T^{-3/7} = \tilde{O}(T^{4/7}),$$

with high probability. Then by taking  $\delta = \delta/T^{3/7}$  we can easily prove that

$$\mathbb{E} \left[ \sum_{t=T_1+1}^T \left\| X_t - \tilde{X}_t \right\|_{V_t^{-1}} \right] = \tilde{O}(T^{4/7}).$$

Similarly, it holds that

$$\mathbb{E} \left[ \sum_{t=T_1+1}^T \|X_t - x_{t,*}\|_{V_t^{-1}} \right] = \tilde{O}(T^{4/7}).$$

Therefore, we have that

$$\sum_{t=T_1+1}^T \left( \mu(x_{t,*}^T \theta) - \mu(\tilde{X}_t^T \theta) \right) = \tilde{O}(T^{4/7}).$$

Finally, the cost of pure exploration is also of scale  $\tilde{O}(T^{4/7})$ , which concludes the proof for UCB-type bandit algorithms.

Second, we will prove Theorem 1 for TS bandit algorithms. Similarly, we would show the detailed proof for the algorithm LinTS, and the regret bound of all other TS algorithms could be similarly deduced:

(1). According to [5], we know that for LinTS (and similarly for other TS algorithms) we have that

$$\mathbb{E} \left[ \sum_{t=T_1+1}^T \left( \mu(x_{t,*}^T \theta) - \mu(\tilde{X}_t^T \theta) \right) \right] = \tilde{O} \left( \sum_{t=T_1+1}^T \left\| \tilde{X}_t \right\|_{V_t^{-1}} + \|x_{t,*}\|_{V_t^{-1}} \right).$$

But for completeness we still offer an alternative proof for this equality:

$$\begin{aligned} \tilde{X}_t^\top \hat{\theta}_t + \alpha^* \left\| \tilde{X}_t \right\|_{V_t^{-1}} Z_t &\geq x_{t,*}^\top \theta + \alpha^* \|x_{t,*}\|_{V_t^{-1}} Z_{t,*} + x_{t,*}^\top (\hat{\theta}_t - \theta) \\ &\geq x_{t,*}^\top \theta + \alpha^* \|x_{t,*}\|_{V_t^{-1}} Z_{t,*} + \|x_{t,*}\|_{V_t^{-1}} \left\| \hat{\theta}_t - \theta \right\|_{V_t} \\ &\geq x_{t,*}^\top \theta + (\alpha^* Z_{t,*} - \alpha(T)) \|x_{t,*}\|_{V_t^{-1}}, \end{aligned}$$

where  $Z_t$  and  $Z_{t,*}$  are IID normal random variables,  $\forall t$ . Therefore, it holds that,

$$\begin{aligned}\tilde{X}_t^\top \theta &\geq x_{t,*}^\top \theta + (\alpha^* Z_{t,*} - \alpha(T)) \|x_{t,*}\|_{V_t^{-1}} - \alpha^* \|X_t\|_{V_t^{-1}} Z_t + X_t^\top (\theta - \hat{\theta}_t), \\ (x_{t,*} - X_t)^\top \theta &\leq (\alpha(T) + \alpha^* Z_t) \|X_t\|_{V_t^{-1}} + (\alpha(T) - \alpha^* Z_{t,*}) \|x_{t,*}\|_{V_t^{-1}} = K_t,\end{aligned}$$

where  $K_t$  is normal random variable with

$$\mathbb{E}(K_t) \leq 2\alpha(T)T^{-1/3}, \quad \text{SD}(K_t) \leq \sqrt{2}\alpha^*T^{-1/3}.$$

Consequently, we have

$$\begin{aligned}\sum_{t=T_1+1}^T (x_{t,*}^\top \theta - \tilde{X}_t^\top \theta) &\leq \sum_{t=T_1+1}^T K_t := K \\ \mathbb{E}(K) &= 2\alpha(T)T^{2/3} = \tilde{O}(T^{2/3}), \quad \text{SD}(K) \leq \sqrt{2}\alpha^*T^{1/6} = O(T^{1/6}).\end{aligned}$$

We have

$$P(K > (2\alpha(T) + \sqrt{2}\alpha^*)T^{2/3}) \leq \frac{1}{c\sqrt{\pi}\sqrt{T}} e^{-c^2 T/2}.$$

This probability upper bound is ultra small and hence negligible. Therefore, we not only prove the expected cumulative regret could be controlled, but also provide a probability bound. Finally, the cost of pure exploration is also of scale  $\tilde{O}(T^{2/3})$ , which concludes the proof.

(2). Although Theorem 1(2). is particularly proposed for UCB-type algorithms, but we would expect a similar result for TS-based algorithms as well. The intuition is presented in Remark 4(2) in detail.

(3). For LinTS, the proof could also be similarly deduced as LinUCB. And we modify the definition of  $\tilde{X}_t$  as

$$\tilde{X}_t = \underset{x \in \mathcal{A}_t}{\operatorname{argmax}} x^\top \hat{\theta}_t + \alpha^* \|x\|_{V_t^{-1}} \tilde{Z}_t,$$

where  $Z_t$  is a standard normal random variable. And we could similarly show that:

$$\begin{aligned}X_t^\top \hat{\theta}_t + \alpha_{i_t} \|X_t\|_{V_t^{-1}} Z_t &\geq x_{t,*}^\top \theta + \alpha_{i_t} \|x_{t,*}\|_{V_t^{-1}} Z_{t,*} + x_{t,*}^\top (\hat{\theta}_t - \theta) \\ &\geq x_{t,*}^\top \theta + \alpha_{i_t} \|x_{t,*}\|_{V_t^{-1}} Z_{t,*} + \|x_{t,*}\|_{V_t^{-1}} \|\hat{\theta}_t - \theta\|_{V_t} \\ &\geq x_{t,*}^\top \theta + (\alpha_{i_t} Z_{t,*} - \alpha(T)) \|x_{t,*}\|_{V_t^{-1}}.\end{aligned}$$

Therefore, it holds that,

$$\begin{aligned}X_t^\top \theta &\geq x_{t,*}^\top \theta + (\alpha_{i_t} Z_{t,*} - \alpha(T)) \|x_{t,*}\|_{V_t^{-1}} - \alpha_{i_t} \|X_t\|_{V_t^{-1}} Z_t + X_t^\top (\theta - \hat{\theta}_t), \\ (x_{t,*} - X_t)^\top \theta &\leq (\alpha(T) + \alpha_{i_t} Z_t) \|X_t\|_{V_t^{-1}} + (\alpha(T) - \alpha_{i_t} Z_{t,*}) \|x_{t,*}\|_{V_t^{-1}} = K_t,\end{aligned}\tag{17}$$

where  $K_t$  is normal random variable with

$$\mathbb{E}(K_t) \leq 2\alpha(T)T^{-2/7}, \quad \text{SD}(K_t) \leq \sqrt{2}\alpha_{i_t} T^{-2/7} \leq \sqrt{2}\alpha^* T^{-2/7}.$$

According to Lemma 5, we have that for arbitrary  $\xi > 0$

$$\begin{aligned}P\left(\max_{t \in T} K_t \geq 2\alpha(T)T^{-2/7} + (\sqrt{2\log(T)} + \xi) \sqrt{2}\alpha^* T^{-2/7}\right) &\leq P\left(\max_{t \in T} \frac{K_t - \mathbb{E}[K_t]}{\text{SD}(K_t)} \geq \sqrt{2\log(T)} + \xi\right) \\ &= T \times P\left(Z \geq \sqrt{2\log(T)} + \xi\right) \quad Z \sim N(0, 1) \\ &\leq T \frac{1}{\sqrt{\pi}(\sqrt{2\log(T)} + \xi)} \exp\left(-(\sqrt{2\log(T)} + \xi)^2/2\right) \\ &\leq \frac{1}{\sqrt{2\log(T)} + \xi} \exp\left(-\frac{\xi^2}{2}\right).\end{aligned}$$

By taking  $\xi = 2\sqrt{\log(T)}$ , it holds that

$$P \left( \max_{t \in T} K_t \geq 2\alpha(T)T^{-2/7} + \left( \sqrt{2\log(T)} + \xi \right) \sqrt{2\alpha^* T^{-2/7}} \right) \leq \frac{1}{2T^2 \sqrt{\log(T)}}.$$

Since this probability upper bound is ultra small and hence negligible, we have

$$(x_{t,*} - X_t)^\top \theta \leq 2\alpha(T)T^{-2/7}, \quad t > T_1.$$

Similarly, we could also deduce that

$$(x_{t,*} - \tilde{X}_t)^\top \theta \leq 2\alpha(T)T^{-2/7}, \quad t > T_1.$$

According to the argument we just made for the proof of Theorem I(3). on LinUCB. (e.g. Eqn. 15), this result implies that

$$\mathbb{E} \left[ \sum_{t=T_1+1}^T \|X_t - x_{t,*}\|_{V_t^{-1}} \right] = \tilde{O}(T^{4/7}), \quad \mathbb{E} \left[ \sum_{t=T_1+1}^T \|\tilde{X}_t\|_{V_t^{-1}} \right] = \tilde{O}(T^{4/7}).$$

According to [5] (or Eqn. (17)), we know that for LinTS we have the similar result as in Eqn. (16):

$$\sum_{t=T_1+1}^T \left( \mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta) \right) = \tilde{O} \left( \sum_{t=T_1+1}^T \|\tilde{X}_t\|_{V_t^{-1}} + \|x_{t,*}\|_{V_t^{-1}} \right).$$

And this directly implies that

$$\sum_{t=T_1+1}^T \left( \mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta) \right) = \tilde{O}(T^{4/7}).$$

Note we could use this procedure to bound the regret for TS bandit algorithms under condition in (3), since most of the proofs for generalized linear bandits are closely related to the rate of  $\sum_{t=T_1+1}^T \|\tilde{X}_t\|_{V_t^{-1}}$ . Finally, the cost of pure exploration is also of scale  $\tilde{O}(T^{4/7})$ , which concludes the proof for Theorem I.  $\square$

### A.3 Analysis of Theorem 2

#### A.3.1 Useful Conclusions

**Proposition 1.** *Assume given the past information  $\mathcal{F}_{t-1}$  and the hyper-parameters to be used by the contextual bandit algorithm at round  $t$ , the arm to be pulled by the contextual bandit algorithm follows a fixed distribution. Denote  $R(\alpha^{(1)}, \dots, \alpha^{(L)}, T, \{\mathcal{F}_{t-1}\})$  as the cumulative regret of the contextual bandit algorithm if it is run with parameters  $(\alpha^{(1)}, \dots, \alpha^{(L)})$  given the past information  $\mathcal{F}_{t-1}$  at round  $t$ . Then the auto tuning method in Algorithm 2 has regret that satisfies the following:*

$$\begin{aligned} \mathbb{E}[R(T)] &\leq \min_{(\alpha^{(1)}, \dots, \alpha^{(L)}) \in J_1 \times \dots \times J_L} \mathbb{E}[R(\alpha^{(1)}, \dots, \alpha^{(L)}, T, \{\mathcal{F}_{t-1}\})] \\ &\quad + 2 \sum_{l=1}^L \sqrt{(e-1)n_l(T-T_1) \log n_l}. \end{aligned}$$

*Proof.* We also reload some notations here for simplicity in the same way as proof of Lemma I in Appendix A.1. More specifically, since at iteration  $t$  we are given the past information  $\mathcal{F}_{t-1}$  to make decision according to different choices of hyper-parameter values, and hence we would omit this notation  $\mathcal{F}_{t-1}$  when we refer to the arm or feature vector we pull under different hyper-parameter values: Denote  $a_t(\alpha_{i_1}^{(1)}, \dots, \alpha_{i_L}^{(L)})$  as the pulled arm at round  $t$  if the hyper-parameters selected at round  $t$  is  $(\alpha_{i_1}^{(1)}, \dots, \alpha_{i_L}^{(L)})$ . Denote  $X_t(\alpha_{i_1}^{(1)}, \dots, \alpha_{i_L}^{(L)})$  as the corresponding feature vector

and  $\mu_t \left( \alpha_{i_1}^{(1)}, \dots, \alpha_{i_L}^{(L)} \right)$  as the corresponding expected reward. It suffices to show that for any  $l = 1, \dots, L$ , the following holds.

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E} \left[ \mu_t \left( \alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(l-1)}^{(l-1)}, \alpha_*^{(l)}, \dots, \alpha_*^{(L)} \right) \right] \\ & - \sum_{t=1}^T \mathbb{E} \left[ \mu_t \left( \alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(l-1)}^{(l-1)}, \alpha_{i_t(l)}^{(l)}, \alpha_*^{(l+1)}, \dots, \alpha_*^{(L)} \right) \right] \\ & \leq 2\sqrt{(e-1)n_l T \log n_l}. \end{aligned} \quad (18)$$

For convenience, we will denote  $\left( \alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(l-1)}^{(l-1)}, \alpha_j^{(l)}, \alpha_*^{(l+1)}, \dots, \alpha_*^{(L)} \right)$  as  $(\alpha_j)$  when there is no ambiguity, which means that the first  $l-1$  hyper-parameters are chosen as  $\alpha_{i_t(s)}^{(s)}$  for  $s = 1, \dots, l-1$ , the  $l$ -th hyper-parameter is chosen with index  $j$  and the rest of the hyper-parameters are chosen as  $\alpha_*^{(s)}$  for  $s = l+1, \dots, L$ . Then the result we want to show in Equation 18 can be written as

$$\sum_{t=1}^T \mathbb{E} \left[ \mu_t \left( \alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(l-1)}^{(l-1)}, \alpha_*^{(l)}, \dots, \alpha_*^{(L)} \right) \right] - \sum_{t=1}^T \mathbb{E} \left[ \mu_t \left( \alpha_{i_t(l)} \right) \right] \leq 2\sqrt{(e-1)n_l T \log n_l}. \quad (19)$$

We will also omit the superscript / subscript  $(l)$  for convenience when there is no ambiguity, so  $p_j^{(l)}(t), w_j^{(l)}(t), \hat{y}_t^{(l)}(j)$  are abbreviated as  $p_j(t), w_j(t), \hat{y}_t(j)$  respectively. Denote  $\mathcal{H}_t = \sigma \left( \alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(l-1)}^{(l-1)}, \alpha_*^{(l+1)}, \dots, \alpha_*^{(L)} \right)$  as the  $\sigma$ -algebra induced by the event that at round  $t$ , the first  $l-1$  hyper-parameters are chosen as  $\alpha_{i_t(s)}^{(s)}$  and for  $s = l+1, \dots, L$ , the hyper-parameters are chosen as  $\alpha_*^{(s)}$ . Given  $\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)$ , denote  $y_t(\alpha_j) = \mu_t(\alpha_j) + \epsilon'$  as the observed reward at round  $t$  if  $\alpha_j$  is chosen as  $\alpha_j^{(l)}$  and the rest hyper-parameters given by  $\mathcal{H}_t$ . Here,  $\epsilon'$  is a hypothetical random noise if arm  $a_t(\alpha_j)$  is pulled at round  $t$ .

Given  $\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)$ , by the above definitions and Algorithm 2,  $\hat{y}_t(j) = y_t(\alpha_j)/p_j(t)$  if  $j = i_t(l)$ . Otherwise,  $\hat{y}_t(j) = 0$ . Since  $p_j(t) \geq \frac{\beta_l}{n_l}$ , we have  $\hat{y}_t(j) \leq \frac{n_l}{\beta_l}$  for all  $j \in [n_l]$  and  $t$ . We also have the following two inequalities.

$$\begin{aligned} & \mathbb{E} \left( \sum_{i=1}^{n_l} p_i(t) \hat{y}_t(i) \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t) \right) = \mathbb{E} (p_{i_t(l)}(t) \hat{y}_t(i_t(l)) \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)) \\ & = \mathbb{E} (y_t(\alpha_{i_t(l)}) \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)) = \mathbb{E} [\mu_t(\alpha_{i_t(l)}) \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)] \end{aligned} \quad (20)$$

$$\begin{aligned} & \mathbb{E} \left( \sum_{i=1}^{n_l} p_i(t) \hat{y}_t(i)^2 \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t) \right) = \mathbb{E} (p_{i_t(l)}(t) \hat{y}_t(i_t(l))^2 \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)) \\ & = \mathbb{E} (y_t(i_t(l)) \hat{y}_t(i_t(l)) \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)) \leq \mathbb{E} (\hat{y}_t(i_t(l)) \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)) \\ & = \mathbb{E} \left( \sum_{i=1}^{n_l} \hat{y}_t(i) \middle| \sigma(\mathcal{F}_{t-1}, \mathcal{H}_t) \right) \end{aligned} \quad (21)$$

For a single  $i \in [n_l]$ , since given  $\mathcal{F}_{t-1}$ ,  $p_i^{(l)}(t)$  is already fixed, which means that the choices of other hyper-parameters do not affect the distribution of  $i_t(l)$ . Moreover,  $a_t(\alpha_i)$  follows a fixed distribution due to the conditions in Theorem 2, i.e., the arm to be pulled follows a fixed distribution given the past information and the hyper-parameters to be used at round  $t$ . Therefore, given  $\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t, a_t(\alpha_i), \epsilon')$ ,  $i = i_t(l)$  is still with probability  $p_i^{(l)}(t)$  for all  $i \in [n_l]$ . So

$$\begin{aligned}
\mathbb{E}(\hat{y}_t(i)|\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)) &= \mathbb{E}[\mathbb{E}(\hat{y}_t(i)|\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t, a_t(\alpha_i), \epsilon'))|\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)] \\
&= \mathbb{E}[y_t(\alpha_i)|\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)] \\
&= \mathbb{E}[\mu_t(\alpha_i)|\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)].
\end{aligned} \tag{22}$$

From Equation 21 and 22, we have

$$\mathbb{E}\left(\sum_{i=1}^{n_l} p_i(t) \hat{y}_t(i)^2 |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right) \leq \mathbb{E}\left(\sum_{i=1}^{n_l} \mu_t(\alpha_i) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right) \tag{23}$$

We still look at the lower bound and upper bound of  $\mathbb{E}[\log \frac{W_{T+1}}{W_1}]$ , but now  $W_t = \sum_{i=1}^{n_l} w_i^{(l)}(t)$ , and we will use the abbreviation  $w_i(t) = w_i^{(l)}(t)$  below for ease of notation.

**Lower bound:**

$$\begin{aligned}
\mathbb{E}\left[\log \frac{w_i(t+1)}{w_i(t)} |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] &= \mathbb{E}\left[\frac{\beta_l}{n_l} \hat{y}_t(i) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] \\
&= \mathbb{E}\left[\frac{\beta_l}{n_l} \mu_t(\alpha_i) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] \quad \text{from Equation 22}
\end{aligned}$$

Take an expectation on both sides and sum over  $t$ , we have

$$\mathbb{E}[\log w_i(T+1)] = \frac{\beta_l}{n_l} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_i)]$$

Therefore, for all  $i \in [n_l]$ ,

$$\mathbb{E}[\log \frac{W_{T+1}}{W_1}] \geq \mathbb{E}[\log w_i(T+1)] - \log n_l = \frac{\beta_l}{n_l} \sum_{t=1}^T \mathbb{E}[\mu_t(\alpha_i)] - \log n_l. \tag{24}$$

**Upper bound:** This part is almost the same as the arguments in Lemma 1, except now that the conditional expectation is taken over  $\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)$ . For completeness, we write out the proof of this part below. Again, we will use  $p_i(t) = p_i^{(l)}(t)$  and  $w_i(t) = w_i^{(l)}(t)$  for convenience.

$$\begin{aligned}
\mathbb{E}\left[\log \frac{W_{t+1}}{W_t} |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] &= \mathbb{E}\left[\log \sum_{i=1}^{n_l} \frac{w_i(t+1)}{W_t} |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] \\
&= \mathbb{E}\left[\log \sum_{i=1}^{n_l} \frac{w_i(t)}{W_t} \exp\left(\frac{\beta_l}{n_l} \hat{y}_t(i)\right) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] \\
&= \mathbb{E}\left[\log \sum_{i=1}^{n_l} \frac{p_i(t) - \frac{\beta_l}{n_l}}{1 - \beta_l} \exp\left(\frac{\beta_l}{n_l} \hat{y}_t(i)\right) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] \quad \text{definition of } p_i(t) \\
&\leq \mathbb{E}\left[\log \sum_{i=1}^{n_l} \frac{p_i(t) - \frac{\beta_l}{n_l}}{1 - \beta_l} \left(1 + \frac{\beta_l}{n_l} \hat{y}_t(i) + \frac{(e-2)\beta_l^2}{n_l^2} \hat{y}_t(i)^2\right) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] \\
&\leq \mathbb{E}\left[\log \left(1 + \sum_{i=1}^{n_l} \left[\frac{\beta_l}{n_l(1-\beta_l)} p_i(t) \hat{y}_t(i) + \frac{(e-2)\beta_l^2}{n_l^2(1-\beta_l)} p_i(t) \hat{y}_t(i)^2\right]\right) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right] \\
&\leq \mathbb{E}\left[\sum_{i=1}^{n_l} \left(\frac{\beta_l}{n_l(1-\beta_l)} p_i(t) \hat{y}_t(i) + \frac{(e-2)\beta_l^2}{n_l^2(1-\beta_l)} p_i(t) \hat{y}_t(i)^2 |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)\right)\right] \\
&\leq \frac{\beta_l}{n_l(1-\beta_l)} \mathbb{E}[\mu_t(\alpha_i) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)] + \frac{(e-2)\beta_l^2}{n_l^2(1-\beta_l)} \sum_{i=1}^{n_l} \mathbb{E}[\mu_t(\alpha_i) |\sigma(\mathcal{F}_{t-1}, \mathcal{H}_t)].
\end{aligned}$$

The first inequality “ $\leq$ ” in the above holds since  $e^x \leq 1 + x + (e - 2)x^2$  for  $x \in [0, 1]$ . Here, we have  $0 \leq \frac{\beta_l}{n_l} \hat{y}_t(i) \leq 1$  because  $p_i(t) \geq \frac{\beta_l}{n_l}$ ,  $0 \leq y_t(\alpha_i) \leq 1$  and  $\hat{y}_t(i) \leq \frac{y_t(\alpha_i)}{p_i(t)}$ . The last inequality is from Equation 20, 23. Take another expectation on both sides, we get

$$\mathbb{E} \left[ \log \frac{W_{t+1}}{W_t} \right] \leq \frac{\beta_l}{n_l(1 - \beta_l)} \mathbb{E} [\mu_t(\alpha_{i_t})] + \frac{(e - 2)\beta_l^2}{n_l^2(1 - \beta_l)} \sum_{i=1}^{n_l} \mathbb{E} [\mu_t(\alpha_i)]$$

By summing the above over  $t$ , we have

$$\mathbb{E} [\log \frac{W_{T+1}}{W_1}] \leq \frac{\beta_l}{n_l(1 - \beta_l)} \sum_{t=1}^T \mathbb{E} [\mu_t(\alpha_{i_t(l)})] + \frac{(e - 2)\beta_l^2}{n_l^2(1 - \beta_l)} \sum_{i=1}^{n_l} \sum_{t=1}^T \mathbb{E} [\mu_t(\alpha_i)] \quad (25)$$

Note that the lower bound in Equation 24 holds for any  $i$ , so it also holds for  $\alpha_*^{(l)}$ . Denote

$$G_{\max} = \sum_{t=1}^T \mathbb{E} \left[ \mu_t(\alpha_{i_t(1)}^{(1)}, \dots, \alpha_{i_t(l-1)}^{(l-1)}, \alpha_*^{(l)}, \dots, \alpha_*^{(L)}) \right].$$

Then

$$\frac{\beta_l}{n_l} G_{\max} - \log n_l \leq \frac{\beta_l}{n_l(1 - \beta_l)} \sum_{t=1}^T \mathbb{E} [\mu_t(\alpha_{i_t(l)})] + \frac{(e - 2)\beta_l^2}{n_l^2(1 - \beta_l)} \sum_{i=1}^{n_l} \sum_{t=1}^T \mathbb{E} [\mu_t(\alpha_i)]$$

We note that  $\sum_{t=1}^T \mathbb{E} [\mu_t(\alpha_i)] \leq T$  for all  $i$ , so

$$\frac{\beta_l}{n_l} G_{\max} - \log n_l \leq \frac{\beta_l}{n_l(1 - \beta_l)} \sum_{t=1}^T \mathbb{E} [\mu_t(\alpha_{i_t(l)})] + \frac{(e - 2)\beta_l^2}{n_l(1 - \beta_l)} T$$

Simplify the above inequality and due to the choice of  $\beta_l$ , we have

$$\begin{aligned} G_{\max} - \sum_{t=1}^T \mathbb{E} [\mu_t(\alpha_{i_t(l)})] &\leq \beta_l G_{\max} + (e - 2)\beta_l T + \frac{(1 - \beta_l)n_l}{\beta_l} \log n_l \\ &\leq 2\sqrt{(e - 1)n_l T \log n_l}. \end{aligned}$$

This concludes the proof of Proposition I.  $\square$

**Lemma 6** (Adapted from Lemma 3). *For any  $\delta < 1$ , under our problem setting in Section 3 with the regularization hyper-parameter  $\lambda \in [\lambda_{\min}, \lambda_{\max}]$  ( $\lambda_{\min} > 0$ ), it holds that for all  $t > 0$ ,*

$$\begin{aligned} \|\hat{\theta}_t - \theta^*\|_{V_t} &\leq \beta_t(\delta), \\ \forall x \in \mathbb{R}^d, |x^\top (\hat{\theta}_t - \theta^*)| &\leq \|x\|_{V_t^{-1}} \beta_t(\delta), \end{aligned}$$

with probability at least  $1 - \delta$ , where

$$\beta_t(\delta) = \sigma \sqrt{\log \left( \frac{(\lambda_{\min} + t)^d}{\delta^2 \lambda_{\min}^d} \right)} + \sqrt{\lambda_{\max}} S.$$

*Proof.* The proof of this Lemma is trivial given Lemma 3. For any  $\lambda \in [\lambda_{\min}, \lambda_{\max}]$ , according to Lemma 3 it holds that, for all  $t > 0$ ,

$$\begin{aligned} \|\hat{\theta}_t - \theta^*\|_{V_t} &\leq \beta_t(\delta), \\ \forall x \in \mathbb{R}^d, |x^\top (\hat{\theta}_t - \theta^*)| &\leq \|x\|_{V_t^{-1}} \beta_t(\delta), \end{aligned}$$

with probability at least  $1 - \delta$ , where

$$\beta_t(\delta) = \sigma \sqrt{\log \left( \frac{(\lambda + t)^d}{\delta^2 \lambda^d} \right)} + \sqrt{\lambda} S \leq \sigma \sqrt{\log \left( \frac{(\lambda_{\min} + t)^d}{\delta^2 \lambda_{\min}^d} \right)} + \sqrt{\lambda_{\max}} S.$$

$\square$

### A.3.2 Proof of Theorem 2

*Proof.* We could validate Theorem 2 by extending the proof of Theorem 1 with Proposition 1. Note that most contextual bandit algorithms contain three types of hyper-parameters: one is the exploration rate, which we have throughout discussed in the proof of Theorem 1. The second class is the stepsize of some gradient-based optimization loop (e.g. Laplace-TS [4]), but the output from the loop when the convergent criteria is met is similar. In other words, this kind of hyper-parameter is not critical in the theoretical proof. The last one is the regularization parameter  $\lambda$ , but it can be easily handled by using Lemma 6. Therefore, we only need to consider the case when we tune the exploration rate and the regularization parameter simultaneously. We will take LinUCB with two hyperparameters (i.e. exploration rate and regularization parameter) as an example:

The proof is similar to the one in Appendix A.2. Denote the candidate sets for hyper-parameter  $\alpha$  and  $\lambda$  as  $J_1$  and  $J_2$  ( $0 < \lambda_{\min} \leq J_2 \leq \lambda_{\max}$ ). And denote  $V_t(\lambda) = \lambda I + \sum_{i=1}^{t-1} X_t X_t^\top$ ,  $\alpha_{i_t}$  and  $\lambda_{i_t}$  as the exploration and regularization rate we tune in our Syndicated framework at round  $t$ . Moreover, we define  $\alpha^* = \min_{\alpha \in J_1} \alpha$ ,  $\lambda^* = \min_{\lambda \in J_2} \lambda$ . With probability at least  $1 - \delta$ ,

$$\left\| \hat{\theta} - \theta \right\|_{V_t(\lambda)} \leq \beta_t(\delta) := \alpha(T, \delta), \quad \forall \lambda \in J_2$$

where the definition of  $\beta_t(\delta)$  is reloaded in Lemma 6. And we would omit  $\delta$  for simplicity. For  $t > T_1$ , we denote the feature vector pulled at round  $t$  as  $X_t$ , i.e.

$$X_t = \underset{x \in \mathcal{A}_t}{\operatorname{argmax}} x^\top \hat{\theta}_t + \alpha_{i_t} \|x\|_{V_t^{-1}(\lambda_{i_t})}, \quad X_t = X_t(\alpha_{i_t}, \lambda_{i_t} | \mathcal{F}_{t-1}).$$

And we also define  $\tilde{X}_t = X_t(\alpha^*, \lambda^* | \mathcal{F}_{t-1})$ , i.e.

$$\tilde{X}_t = \underset{x \in \mathcal{A}_t}{\operatorname{argmax}} x^\top \hat{\theta}_t + \alpha^* \|x\|_{V_t^{-1}(\lambda^*)}.$$

According to Proposition 1, it holds that

$$\begin{aligned} \mathbb{E}[R(T)] &\leq \mathbb{E}[R(\alpha^*, \lambda^*, T, \{\mathcal{F}_{t-1}\})] + O(\sqrt{T - T_1}) \\ &\leq \mathbb{E} \left[ \sum_{t=T_1+1}^T \left( \mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta) \right) \right] + O(\sqrt{T - T_1}) \end{aligned}$$

According to the proof of LinUCB we could similarly argue that

$$\begin{aligned} x_{t,*}^\top \theta - \tilde{X}_t^\top \theta &\leq \alpha^* \left( \left\| \tilde{X}_t \right\|_{V_t^{-1}(\lambda^*)} - \left\| \tilde{x}_{t,*} \right\|_{V_t^{-1}(\lambda^*)} \right) + \left\| x_{t,*} - \tilde{X}_t \right\|_{V_t^{-1}(\lambda^*)} \left\| \hat{\theta}_t - \theta \right\|_{V_t(\lambda^*)} \\ &\leq (\alpha^* + \alpha(T)) \left\| \tilde{X}_t \right\|_{V_t^{-1}(\lambda^*)} + \alpha(T) \left\| x_{t,*} \right\|_{V_t^{-1}(\lambda^*)}. \end{aligned}$$

In conclusion, we have that

$$\sum_{t=T_1+1}^T \left( \mu(x_{t,*}^\top \theta) - \mu(\tilde{X}_t^\top \theta) \right) = \tilde{O} \left( \sum_{t=T_1+1}^T \left\| \tilde{X}_t \right\|_{V_t^{-1}(\lambda^*)} + \sum_{t=T_1+1}^T \left\| x_{t,*} \right\|_{V_t^{-1}(\lambda^*)} \right).$$

By Lemma 4 and choosing  $T_1 = T^{2/3}$ , it holds that,

$$\sum_{t=T_1+1}^T \left\| x_{t,*} \right\|_{V_t^{-1}}, \quad \sum_{t=T_1+1}^T \left\| \tilde{X}_t \right\|_{V_t^{-1}} = O(T \times T^{-1/3}) = O(T^{2/3}).$$

Note we could literally use the identical argument for all UCB and TS bandit algorithms as in the proof of Theorem 1, and the only modification is the value of  $\alpha(T)$  we newly defined in Lemma 6.

To prove the Theorem 1(3) holds, we can also use an exactly identical argument as in the proof of Theorem 1(3) in Appendix A.2, and the only difference is we replace the value of  $\alpha(T)$  in our main paper by the newly defined one in Lemma 6, and hence we would not copy it here again. And this fact concludes our proof.  $\square$

#### A.4 Experimental Settings

**Simulations.** We use  $d = 10$ ,  $K = 100$  and draw  $\theta^* \sim \text{Uniform}(-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}})$ . For linear bandits, we draw the feature vectors  $x_{t,a} \sim \text{Uniform}(-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}})$  and transform the mean reward of arm  $a$  at round  $t$  by  $\mu_{t,a} \leftarrow \frac{\mu_{t,a} + 1}{2}$  to make sure the mean rewards are within  $[0, 1]$ . Each round an arm is pulled, a sample reward  $Y_t \sim N(\mu_{t,a}, 0.1)$  is revealed to the player. For logistic models, the feature vectors  $x_{t,a} \sim \text{Uniform}(-1, 1)$  and the corresponding mean reward is  $\mu_{t,a} = 1/(1 + \exp(-x_{t,a}^T \theta^*))$ . A sample Bernoulli reward is drawn when an arm is pulled.

**Real datasets.** We use the benchmark MovieLens 100K dataset similarly as in [8]. The MovieLens dataset contains 100K ratings on 1,682 movies contributed by 943 users. For data preprocessing, we apply LIBPMF [26, 27] to factorize the ratings matrix to get feature matrices for both users and movies with  $d = 20$ . We randomly select  $K = 1000$  movies (arms) in each round, and the model parameter  $\theta^*$  is defined as the averaged feature vectors of 100 randomly selected users. For linear models, the mean reward is defined as  $\mu_{t,a} = x_{t,a}^T \theta$  and transformed into  $[0, 1]$ . The sample reward is drawn from  $N(\mu_{t,a}, 1)$ . For logistic models, the mean reward is defined as  $\mu_{t,a} = 1/(1 + \exp(-x_{t,a}^T \theta^*))$ , and the sample reward is drawn from a Bernoulli distribution.

#### A.5 Additional experiments on tuning SGD-TS

In this section, we show the comparison of different tuning methods in SGD-TS [14], a recently proposed efficient algorithm for generalized linear bandit. We apply SGD-TS with a logistic model to the datasets considered in Section 6. SGD-TS has four tuning parameters, the length of epoch  $\tau$ , two exploration parameters  $\alpha^{(1)}$  and  $\alpha^{(2)}$ , step size for stochastic gradient descent  $\eta_0$ . In [14], the experiments are conducted by using a grid search of all four parameters, which is not feasible in practice. Since the epoch length has to be pre-determined, it is not applicable to tune it online. We set  $\tau = 10 \times \lfloor \max(\log T, d) \rfloor$  as suggested by the grid search set in [14] and fix it for all tuning methods. The tuning set for  $\alpha^{(1)}$  and  $\alpha^{(2)}$  are the same  $\{0, 0.01, 0.1, 1, 10\}$ . The tuning set for step size  $\eta_0$  is set as  $\{0.01, 0.1, 1, 10\}$ . The theoretical choices of step size  $\eta_0$  in SGD-TS are intractable, so for the tuning methods in Section 6, we make the following modifications:

1. **OP** [9]: We modify OPLINUCB to tune step size  $\eta_0$  only.
2. **Corral** [3]: We modify the CORRAL model selection framework to tune step size  $\eta_0$  only.
3. **Corral-Combined** [3]: We modify the CORRAL model selection framework to tune all three hyper-parameters  $\alpha^{(1)}$ ,  $\alpha^{(2)}$  and  $\eta_0$ . And the tuning set contain all possible combinations of these three hyper-parameters.
4. **TL (Our work, Algorithm 1)**: This is our proposed Algorithm 1, where we use the two-layer bandit structure to tune the step size  $\eta_0$  only.
5. **TL-Combined (Our work, Algorithm 1)**: This method tunes all three hyper-parameters  $\alpha^{(1)}$ ,  $\alpha^{(2)}$  and  $\eta_0$  using Algorithm 1, but with the tuning set containing all the possible combinations of the three hyper-parameters.
6. **Syndicated (Our work, Algorithm 2)**: This method keeps three separate tuning sets for  $\alpha^{(1)}$ ,  $\alpha^{(2)}$  and  $\eta_0$  respectively. It uses the Syndicated Bandits framework in Algorithm 2.

For OP Corral and TL, since they do not tune the two exploration parameters,  $\alpha^{(1)}$  and  $\alpha^{(2)}$  are set as the theoretical values as in [14]. Results reported in Figure 4 are averaged over 10 repeated experiments. From the plots, we can see that 1) our proposed Syndicated Bandits framework outperforms TL-combined method since now there are in total three hyper-parameters and the regret of TL-combined depends on the number of hyper-parameters exponentially. 2) Tuning all 3 hyper-parameters significantly outperforms tuning only the step size as in OP, Corral and TL. This further indicates that tuning multiple hyper-parameters is better than tuning fewer. On the other hand, it suggests that the theoretical choices of the exploration parameters do not always perform better than the fine-tuned results. 3) Our proposed TL algorithm outperforms OP and Corral when tuning only the step size.

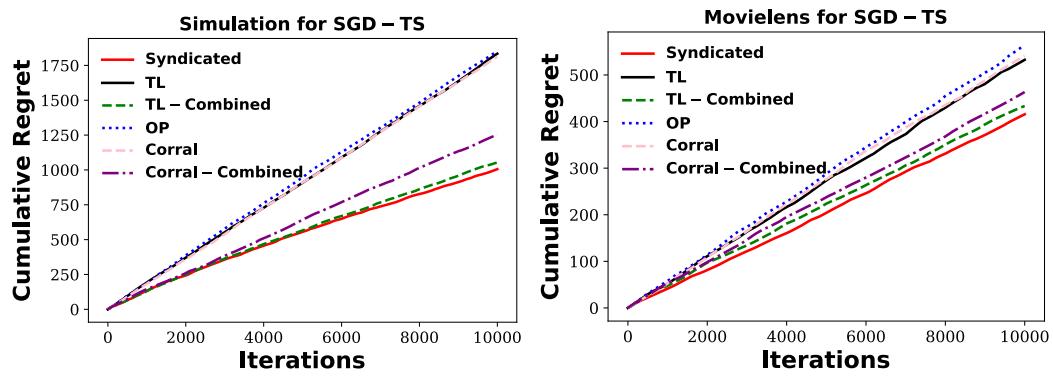


Figure 4: Comparison of hyper-parameters selection methods in SGD-TS.