

# An Information-Theoretic Analysis of Thompson Sampling with Infinite Action Spaces

Amaury Gouverneur, Borja Rodríguez-Gálvez, Tobias J. Oechtering, and Mikael Skoglund

Division of Information Science and Engineering (ISE)

KTH Royal Institute of Technology

{amauryg, borjarg, oech, skoglund}@kth.se

**Abstract**—This paper studies the Bayesian regret of the Thompson Sampling algorithm for bandit problems, building on the information-theoretic framework introduced by Russo and Van Roy [1]. Specifically, it extends the rate-distortion analysis of Dong and Van Roy [2], which provides near-optimal bounds for linear bandits. A key limitation of these results is the assumption of a finite action space. We address this by extending the analysis to settings with infinite and continuous action spaces. Additionally, we specialize our results to bandit problems with expected rewards that are Lipschitz continuous with respect to the action space, deriving a regret bound that explicitly accounts for the complexity of the action space.

**Index Terms**—multi-armed bandit, logistic bandit, Thompson Sampling, information theory, regret bounds, online optimization.

## I. INTRODUCTION

Bandit problems are a class of sequential decision-making problems where agents interact with an unknown environment by selecting actions and receiving rewards in return. The agent’s objective is to maximize its expected cumulative reward, which is the expected sum of rewards obtained during its interaction with the environment. This requires balancing exploration—trying different actions to gather information for future rewards—with exploitation, where the agent leverages known actions to maximize its immediate gain. Bandit problems have attracted significant attention due to their wide range of applications in fields such as healthcare, finance, recommender systems, and telecommunications (see [3] for a survey on various applications). The theoretical evaluation of an algorithm’s performance in bandit problems is typically done by studying the *expected regret*, which measures the gap between the cumulative reward earned by the algorithm and the reward an oracle would achieve by always selecting the optimal action [4]–[11]. One effective method for achieving small regret is the Thompson Sampling (TS) algorithm [12], which, despite its simplicity, has demonstrated remarkable performance [5, 13, 14].

Studying the Thompson Sampling regret, Russo and Van Roy [1] introduced the concept of *information ratio*. This statistic captures the trade-off between achieving immediate low regret and gaining information about the optimal action. Using this concept, they provide a general upper bound that

depends on the prior entropy of the optimal action  $H(A^*)$ . A limitation of this result is that the prior entropy of the optimal action can grow arbitrarily large with the number of actions or get infinite if the action space is continuous. Combining the previous techniques with a rate-distortion analysis, Dong and Van Roy [2] were able to control the regret of Thompson Sampling via the entropy of a *statistic* of the parameter  $\Theta$ . This approach provides Bayesian regret bounds that remain sharp with large action spaces. In particular, they derived a near-optimal regret rate of  $O(d\sqrt{T\log T})$  for  $d$ -dimensional linear bandit problems. However, their results are limited to bandit problems with finite action and environment space.

In this work, we extend the results of [2] to settings with infinite and continuous action and parameter spaces. For bandit problems where the expected rewards are Lipschitz continuous with respect to a metric action space, we derive a regret bound that explicitly depends on the complexity of the action space.

## II. PRELIMINARIES

### A. General Notation

We denote random variables by capital letters (e.g.,  $X$ ), their specific realizations by lowercase letters (e.g.,  $x$ ), and their outcome spaces using calligraphic notation (e.g.,  $\mathcal{X}$ ). The distribution of a random variable  $X$  is represented as  $\mathbb{P}_X$ .

### B. Problem setup

We consider a sequential decision-making problem where, at each time step  $t \in 1, \dots, T$ , an agent interacts with the environment by selecting an action  $A_t \in \mathcal{A}$  from a set of available actions  $\mathcal{A}$ . Based on the chosen action, the agent receives a real-valued reward  $R_t \in \mathbb{R}$ . The action-reward pair  $(A_t, R_t)$  is then added to the history, updating it to  $H^{t+1} = H^t \cup \{A_t, R_t\}$ , which will be available for decision-making in the next round. This process continues until  $t = T$ .

In the Bayesian framework, the environment is characterized by a parameter  $\theta \in \mathcal{O}$ , which is unknown to the agent and drawn from a known prior distribution  $\mathbb{P}_\Theta$ . Together with the reward distribution  $\mathbb{P}_{R|A,\Theta}$ , this prior fully defines the bandit problem. Since the reward depends on the chosen action and the parameters, it can be expressed as  $R_t = R(A_t, \Theta)$  for some possibly random function  $R : \mathcal{A} \times \mathcal{O} \rightarrow \mathbb{R}$ .

The goal of the agent is to select actions that maximize the total accumulated reward. Specifically, the agent aims to learn a policy  $\pi = \{\pi_t : \mathcal{H}^t \rightarrow \mathcal{A}\}_{t=1}^T$  that, at each time  $t \in \{1, \dots, T\}$ , selects an action  $A_t$  based on the history  $H^t$ . The objective is to find a policy that maximize the *expected cumulative reward*  $R_T(\pi) := \mathbb{E} \left[ \sum_{t=1}^T R(\pi_t(H^t), \Theta) \right]$ .

If the agent had access to the parameter  $\Theta$ , it could always select the action that maximizes the expected reward defined by the mapping  $\pi_* : \mathcal{O} \rightarrow \mathcal{A}$  as  $\pi_*(\theta) = \operatorname{argmax}_{a \in \mathcal{A}} \mathbb{E}[R(a, \theta)]$ . We refer to such an action as the *optimal action*  $A^* := \pi_*(\Theta)$ . To ensure such a policy exists, we make the technical assumption that the set of actions  $\mathcal{A}$  is compact. Following [15], we assume that there exists a mapping  $\pi_*$  is one-to-one<sup>1</sup>.

We evaluate the performance of a policy  $\pi$  with its Bayesian regret, defined as:

$$\mathbb{E}[\text{Regret}(T)] := \mathbb{E} \left[ \sum_{t=1}^T R(A^*, \Theta) - R(A_t, \Theta) \right],$$

where the actions  $A_t$  are selected by the policy  $\pi$ , and the expectation is taken over the randomness of the action selection, the reward distribution, and the distribution of  $\Theta$ .

### C. Thompson Sampling algorithm

Thompson Sampling is an efficient algorithm for solving bandit problems. It works by selecting actions randomly based on their posterior probability of being optimal. Specifically, at each time step  $t \in \{1, \dots, T\}$ , the agent samples a parameter estimate  $\hat{\Theta}_t$  from the posterior distribution, conditioned on the history  $H^t$ . The agent then selects the action that is optimal for the sampled parameter,  $A_t = \pi_*(\hat{\Theta}_t)$ , receives a reward  $R_t$ , and updates the history to  $H^{t+1} = H^t \cup \{\hat{A}_t, R_t\}$ . The pseudocode for Thompson Sampling is given in Algorithm 1.

---

#### Algorithm 1 Thompson Sampling algorithm

---

- 1: **Input:** parameter prior  $\mathbb{P}_\Theta$ .
  - 2: **for**  $t = 1$  **to**  $T$  **do**
  - 3:   Sample a parameter estimate  $\hat{\Theta}_t \sim \mathbb{P}_{\Theta|H^t}$ .
  - 4:   Take the corresponding optimal action  $A_t = \pi_*(\hat{\Theta}_t)$ .
  - 5:   Collect the reward  $R_t = R(A_t, \Theta)$ .
  - 6:   Update the history  $H^{t+1} = H^t \cup \{A_t, R_t\}$ .
  - 7: **end for**
- 

### D. Notations Specific to Bandit Problems

Since the  $\sigma$ -algebras of the history are often used in conditioning, we introduce the notations  $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot|H^t]$  and  $\mathbb{P}_t[\cdot] := \mathbb{P}[\cdot|H^t]$  to denote the conditional expectation and probability given  $H^t$ . Additionally, we define  $I_t(A^*; R_t|A_t) :=$

$\mathbb{E}_t[\text{D}_{\text{KL}}(\mathbb{P}_{R_t|H^t, A^*, A_t} \| \mathbb{P}_{R_t|H^t, A_t})]$  as the disintegrated conditional mutual information between the optimal action  $A^*$  and the reward  $R_t$  conditioned on action  $A_t$ , given the history  $H^t$ .

### III. INFORMATION RATIO AND COMPRESSED TS

In their analysis of Thompson Sampling for bandit problems, Russo and Van Roy [1] introduced an important quantity, the *information ratio* defined as the following random variable:

$$\Gamma_t(A^*, A_t) := \frac{\mathbb{E}_t[R(A^*, \Theta) - R(A_t, \Theta)]^2}{I_t(A^*; R(A_t, \Theta), A_t)}.$$

This ratio measures the trade-off between minimizing the current squared regret and gathering information about the optimal action. In other words, a small information ratio implies that any significant regret has been compensated by a substantial gain in information about the optimal action.

Russo and Van Roy use this concept to provide a general regret bound of  $\sqrt{\Gamma \cdot T \cdot H(A^*)}$ , which depends on the time horizon  $T$ , the entropy of the prior distribution of  $A^*$ , and an algorithm- and problem-dependent upper bound  $\Gamma$  on the average expected information ratio [1, Proposition 1]. A limitation of this approach is that the prior entropy of the optimal action,  $H(A^*)$ , can grow arbitrarily large with the number of actions or get infinite if the action space is continuous. Dong and Van Roy [2] extended this analysis by controlling the regret of the Thompson Sampling through the regret of an approximate learning they refer to as *one-step compressed Thompson Sampling*. Under a continuity assumption of the expected reward with respect to the action space, they upper bound the regret via a *compressed statistic*  $\Theta_\varepsilon$  of the parameter  $\Theta$ , along with the information ratio of the one-step compressed Thompson Sampling defined as

$$\tilde{\Gamma}_t(\tilde{\Theta}_t^*, \tilde{\Theta}_t) := \frac{\mathbb{E}_t[R(\pi_*(\tilde{\Theta}_t^*), \Theta) - R(\pi_*(\tilde{\Theta}_t), \Theta)]^2}{I_t(\tilde{\Theta}_t^*; R(\pi_*(\tilde{\Theta}_t), \Theta), \tilde{\Theta}_t)}.$$

for some carefully crafted variables  $\tilde{\Theta}_t^*$  and  $\tilde{\Theta}_t$  which depend respectively on  $\Theta_\varepsilon$  and  $\hat{\Theta}_{\varepsilon, t}$  (a statistic corresponding to the parameter estimate  $\hat{\Theta}_t$ ). The resulting bound is of the form  $\varepsilon \cdot T + \sqrt{\tilde{\Gamma} \cdot T \cdot H(\Theta_\varepsilon)}$  where  $\tilde{\Gamma}$  is a problem dependent upper bound on the average expected  $\tilde{\Gamma}_t(\tilde{\Theta}_t^*, \tilde{\Theta}_t)$  and  $\varepsilon$  is a parameter that controls the quantization of  $\Theta_\varepsilon$ . For  $d$ -dimensional linear bandit, they show  $\tilde{\Gamma} \leq d/2$  in [2, Proposition 3] and after optimizing over the choice of  $\varepsilon$  they derive a near-optimal regret rate of  $O(d\sqrt{T} \log T)$ .

However, those results do not hold for infinite or continuous action and parameter spaces, as the proof techniques for constructing the one-step compressed Thompson Sampling variables  $\tilde{\Theta}_t^*$  and  $\tilde{\Theta}_t$  [2, Proposition 2] specifically require a finite parameter space or finite action support. We address this limitation in Section IV, and, to simplify the exposure of the results, we adapt their construction such that it depends on a *statistic*  $A_\varepsilon^*$  of the optimal action  $A^*$  instead of a statistic  $\Theta_\varepsilon$  of the parameter  $\Theta$ .

<sup>1</sup>Note that Thompson Sampling disregards actions that are not optimal for any parameter. If a particular action is optimal for several parameters, we can include duplicate versions in the action set to ensure a one-to-one correspondence between each parameter and its optimal action (and similarly if a particular parameter is optimal for several actions).

#### IV. MAIN RESULTS

In this section, we explain how we extend the results of Dong and Van Roy [2] to continuous and infinite parameter spaces. We then apply the regret bound to linear bandits. To simplify the expressions, we will use the notation  $R(A_t)$  as a shorthand for  $R(A_t, \Theta)$ .

We begin by adapting their construction of a one-step compressed Thompson Sampling such that it depends on a statistic  $A_\varepsilon^*$  of the optimal action  $A^*$  and a corresponding carefully crafted action sampling function  $\phi_t : \mathcal{A} \rightarrow \mathcal{A}$ , such that  $\tilde{A}_t^* := \phi_t(A_\varepsilon^*)$ , for each round  $t \in \{1, \dots, T\}$ . Similar to [2, Proposition 2], this statistic and sampling functions are designed to simultaneously satisfy the following three requirements:

- (i) The statistic  $A_\varepsilon^*$  is less informative than  $A^*$ , that is,  $H(A_\varepsilon^*) \leq H(A^*)$ .
- (ii) At each round  $t \in \{1, \dots, T\}$ , the one-step compressed Thompson Sampling regret is “ $\varepsilon$ -close” to the Thompson Sampling regret.
- (iii) For each time step  $t \in \{1, \dots, T\}$ , the one-step compressed Thompson Sampling regret can be bounded using the information gained about the statistic  $A_\varepsilon^*$ . At the same time, it reveals no more information about  $A_\varepsilon^*$  than Thompson Sampling.

Following [2], we construct a partition of  $\{\mathcal{A}_k\}_{k=1}^K$  of  $\mathcal{A}$  such that for each  $k = 1, \dots, K$ , for all  $a, a' \in \mathcal{A}_k$ ,  $\mathbb{E}[R(a, \pi_\star^{-1}(a)) - R(a', \pi_\star^{-1}(a))] \leq \varepsilon$ . We define the statistic  $A_\varepsilon^*$  as the random variable recording the partition of  $A^*$ :

$$A_\varepsilon^* = k \iff A^* \in \mathcal{A}_k. \quad (1)$$

To prove the existence of a one-step compressed Thompson Sampling satisfying requirements (i), (ii) and (iii), we introduce Proposition 1.

**Proposition 1.** *Consider a space  $\mathcal{A}$ , two functions  $f : \mathcal{A} \rightarrow \mathbb{R}_+$  and  $g : \mathcal{A} \rightarrow \mathbb{R}_+$ , and a probability distribution  $\mathbb{Q}$  on  $\mathcal{A}$ . Then, there exists a pair  $(a_1, a_2) \in \mathcal{A}^2$  and a  $q \in [0, 1]$  such that*

$$qf(a_1) + (1-q)f(a_2) \leq \int_{a \in \mathcal{A}} f(a) d\mathbb{Q}(a)$$

and

$$qg(a_1) + (1-q)g(a_2) \leq \int_{a \in \mathcal{A}} g(a) d\mathbb{Q}(a).$$

*Proof.* Let  $\bar{F} = \int_{a \in \mathcal{A}} f(a) d\mathbb{Q}(a)$  and  $\bar{G} = \int_{a \in \mathcal{A}} g(a) d\mathbb{Q}(a)$ . Now, consider the spaces  $\mathcal{A}_f := \{a \in \mathcal{A} : f(a) \leq \bar{F}\}$  and  $\mathcal{A}_g := \{a \in \mathcal{A} : g(a) \leq \bar{G}\}$ . If  $\mathcal{A}_f \cap \mathcal{A}_g \neq \emptyset$ , then taking both  $a_1$  and  $a_2$  from  $\mathcal{A}_f \cap \mathcal{A}_g$  trivially satisfies the conditions for all  $q \in [0, 1]$ . Therefore, let us assume that the sets are disjoint for the rest of the proof.

Consider some  $a_1 \in \mathcal{A}_f = \mathcal{A}_f^c$  and some  $a_2 \in \mathcal{A}_g = \mathcal{A}_g^c$ . We can rewrite the required condition from the lemma as

$$q \geq \frac{f(a_2) - \bar{F}}{f(a_2) - f(a_1)} \quad \text{and} \quad q \leq \frac{\bar{G} - g(a_2)}{g(a_1) - g(a_2)},$$

where the first inequality took into account that  $f(a_1) < f(a_2)$  by the definition of the sets  $\mathcal{A}_f$  and  $\mathcal{A}_g = \mathcal{A}_f^c$ . This inequality can, in turn, be written as

$$\frac{f(a_2) - \bar{F}}{f(a_2) - f(a_1)} \leq \frac{\bar{G} - g(a_2)}{g(a_1) - g(a_2)}$$

which is equivalent to

$$f(a_2)g(a_1) - \bar{F}(g(a_1) - g(a_2)) \leq \bar{G}(f(a_2) - f(a_1)) + f(a_1)g(a_2).$$

At this point, we have all the ingredients to prove the statement by contradiction. Assume that there is no pair  $(a_1, a_2) \in \mathcal{A}_f \times \mathcal{A}_g$  such that the condition holds, then

$$f(a_2)g(a_1) - \bar{F}(g(a_1) - g(a_2)) > \bar{G}(f(a_2) - f(a_1)) + f(a_1)g(a_2)$$

for every pair  $(a_1, a_2) \in \mathcal{A}_f \times \mathcal{A}_g$ . Therefore, we can integrate over all such pairs, and the inequality should still hold, namely

$$\begin{aligned} & \int_{\mathcal{A}_f} \int_{\mathcal{A}_g} \left[ f(a_2)g(a_1) - \bar{F}(g(a_1) - g(a_2)) \right] d\mathbb{Q}(a_1) d\mathbb{Q}(a_2) \\ & > \int_{\mathcal{A}_f} \int_{\mathcal{A}_g} \left[ \bar{G}(f(a_2) - f(a_1)) + f(a_1)g(a_2) \right] d\mathbb{Q}(a_1) d\mathbb{Q}(a_2). \end{aligned}$$

We must introduce some notation to show that the above inequality cannot happen. Let  $F^- := \int_{\mathcal{A}_f} f(a) d\mathbb{Q}(a)$  and  $F^+ := \int_{\mathcal{A}_g} f(a) d\mathbb{Q}(a)$  and note that  $F^+ + F^- = \bar{F}$ . Similarly,  $G^- := \int_{\mathcal{A}_g} g(a) d\mathbb{Q}(a)$  and  $G^+ := \int_{\mathcal{A}_f} g(a) d\mathbb{Q}(a)$  and  $G^+ + G^- = \bar{G}$ . Using this notation, we can use Fubini's theorem and rewrite the above inequality as

$$F^+G^+ - (F^+ + F^-)(G^+ - G^-) > (G^+ + G^-)(F^+ - F^-) + F^-G^-$$

which can be simplified to  $F^-G^- > F^+G^+$  and which is impossible by the definition of  $F^-$ ,  $F^+$ ,  $G^+$  and  $G^-$ , completing the contradiction and therefore the proof.  $\square$

Equipped with Proposition 1, we can now extend [2, Proposition 2] to infinite and continuous action space.

**Proposition 2.** *Let  $A_\varepsilon^*$  be defined as in (1). For each time step  $t \in \{1, \dots, T\}$ , there exists a of random functions  $\phi_t$  that satisfies the following:*

- 1)  $\mathbb{E}_t[R(A^*) - R(A_t)] \leq \varepsilon + \mathbb{E}_t[R(\tilde{A}_t^*) - R(\tilde{A}_{t,\varepsilon})]$ .
- 2)  $\mathbb{I}_t(A_\varepsilon^*; \tilde{A}_{t,\varepsilon}, R(\tilde{A}_{t,\varepsilon})) \leq \mathbb{I}_t(A_\varepsilon^*; \hat{A}_t, R(\hat{A}_t))$ .

where we set  $\tilde{A}_t^* = \phi_t(A_\varepsilon^*)$  and  $\tilde{A}_{t,\varepsilon} = \phi_t(\hat{A}_{t,\varepsilon})$ , and where  $\hat{A}_{t,\varepsilon}$  is the statistic corresponding to  $\hat{A}_t$ , thus  $A_\varepsilon^*$  and  $\hat{A}_{t,\varepsilon}$  are identically distributed conditioned on  $H^t$ .

*Proof.* The proof follows the technique from [2, Proof of Proposition 2] closely with the key difference that we use Proposition 1 instead of [2, Lemma 1].  $\square$

Adjusting to our construction of *one-step compressed TS* based on  $A_\varepsilon^*$ , we adapt the definition of information ratio as

$$\tilde{\Gamma}_t(\tilde{A}_t^*, \tilde{A}_{t,\varepsilon}) := \frac{\mathbb{E}_t[R(\tilde{A}_t^*, \Theta) - R(\tilde{A}_{t,\varepsilon}, \Theta)]^2}{\mathbb{I}_t(A_\varepsilon^*; \tilde{A}_{t,\varepsilon}, R(\tilde{A}_{t,\varepsilon}))}$$

We can now state and prove our main theorem.

**Theorem 1.** Let  $A_\varepsilon^*$  be defined as in (1) and let  $\tilde{A}_t^*$  and  $\tilde{A}_{t,\varepsilon}$  satisfy Proposition 1. Assume that the average expected one-step compressed TS information ratio is bounded,  $\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\tilde{\Gamma}_t(\tilde{A}_t^*, \tilde{A}_{t,\varepsilon})] \leq \tilde{\Gamma}$ , for some  $\tilde{\Gamma} > 0$ . Then, the TS cumulative regret is bounded as

$$\mathbb{E}[\text{Regret}(T)] \leq \sqrt{\tilde{\Gamma}TH(A_\varepsilon^*)} + \varepsilon T.$$

*Proof.* The proof follows the techniques of [2, Proof of Theorem 1]. We start the proof by upper bounding the regret of TS by the regret of the one-step compressed TS using the first inequality in Proposition 2:

$$\mathbb{E}[\text{Regret}(T)] \leq \varepsilon T + \sum_{t=1}^T \mathbb{E}[R(\tilde{A}_t^*) - R(\tilde{A}_{t,\varepsilon})]. \quad (2)$$

We rewrite the expected regret of one-step compressed TS using the definition of the information ratio:

$$\mathbb{E}[R(\tilde{A}_t^*) - R(\tilde{A}_{t,\varepsilon})] = \mathbb{E} \left[ \sqrt{\tilde{\Gamma}_t(\tilde{A}_t^*, \tilde{A}_{t,\varepsilon}) I_t(A_\varepsilon^*; \tilde{A}_{t,\varepsilon}, R(\tilde{A}_{t,\varepsilon}))} \right]$$

We continue using Jensen's inequality and applying Cauchy-Schwartz inequality:

$$(2) \leq \varepsilon T + \sqrt{\tilde{\Gamma}T \sum_{t=1}^T I(A_\varepsilon^*; \tilde{A}_{t,\varepsilon}, R(\tilde{A}_{t,\varepsilon}) | H^t)}$$

where in the last inequality, we used the assumption that  $\sum_{t=1}^T \mathbb{E}_t[\tilde{\Gamma}_t(\tilde{A}_t^*, \tilde{A}_{t,\varepsilon})] \leq \tilde{\Gamma}T$ . Using the second inequality in Proposition 2, summing the  $T$  mutual information  $I(\Theta_\varepsilon; R_t | H^t, A_t)$  and applying the chain rule, we obtain

$$\mathbb{E}[\text{Regret}(T)] \leq \varepsilon T + \sqrt{\tilde{\Gamma}TI(A_\varepsilon^*; H^T)}.$$

Finally, we upper bound the mutual information  $I(A_\varepsilon^*; H^T)$  by the entropy  $H(A_\varepsilon^*)$  to obtain the claimed result.  $\square$

For bandit problems with expected rewards that are  $L$ -Lipschitz with respect to a metric action space  $(\mathcal{A}, \rho)$ , we can derive a regret bound from Theorem 1 that depends on a measure of the complexity of the action space.

**Corollary 1.** Assume that  $\mathbb{E}[R(a, \theta)]$  is  $L$ -Lipschitz with respect to the action space  $(\mathcal{A}, \rho)$  for all  $\theta \in \mathcal{O}$ , and assume that the average expected one-step compressed TS information ratio is bounded,  $\frac{1}{T} \sum_{t=1}^T \mathbb{E}[\tilde{\Gamma}_t(\tilde{A}_t^*, \tilde{A}_{t,\varepsilon})] \leq \tilde{\Gamma}$ , for some  $\tilde{\Gamma} > 0$ . Then, the TS cumulative regret is bounded as

$$\mathbb{E}[\text{Regret}(T)] \leq \sqrt{\tilde{\Gamma}T \log(\mathcal{N}(\mathcal{A}, \rho, \varepsilon))} + L\varepsilon T,$$

where  $\mathcal{N}(\mathcal{A}, \rho, \varepsilon)$  is the covering number for the metric space  $(\mathcal{A}, \rho)$  at scale  $\varepsilon \geq 0$ .

*Proof.* Under the Lipschitz continuity property of the expected rewards, we note that setting  $A_\varepsilon^*$  as the quantization of  $A^*$  on a  $\varepsilon/L$ -net for  $(\mathcal{A}, \rho)$  satisfies (1). Starting from Theorem 1, we upper bound the entropy  $H(A_\varepsilon^*)$  by the logarithm of the cardinality of  $A_\varepsilon^*$ . Choosing the  $\varepsilon/L$ -net with the smallest cardinality and setting  $\varepsilon = \varepsilon/L$  yields the claimed result.  $\square$

## V. APPLICATION TO LINEAR BANDIT PROBLEMS

In linear bandits problems, both the actions and the environment parameter are parameterized by a feature vector, and the associated expected reward can be written as their inner product. Mathematically, a  $d$ -dimensional linear bandit problem is a bandit problem with  $\mathcal{A}, \mathcal{O} \subset \mathbb{R}^d$  and such that for all  $a \in \mathcal{A}$  and all  $\theta \in \mathcal{O}$  we have

$$\mathbb{E}[R(a, \theta)] = \langle a, \theta \rangle,$$

where the expectation is taken over the reward distribution.

Adapting Corollary 1 to linear bandits, we get the following.

**Proposition 3.** For  $d$ -linear bandit problems with rewards bounded in  $[-1, 1]$ , the TS cumulative regret is bounded as

$$\mathbb{E}[\text{Regret}(T)] \leq \sqrt{2dT \log(\mathcal{N}(\mathcal{A}, \|\cdot\|_2, \varepsilon))} + \varepsilon T$$

*Proof.* From [2, Proposition 3], we have that the average expected one-step compressed TS information ratio is bounded by  $2d$ . Then, noting that linear bandits are 1-Lipschitz with respect to  $(\mathcal{A}, \|\cdot\|_2)$  where  $\|\cdot\|_2$  is the Euclidean distance, and applying Corollary 1 gives the claimed result.  $\square$

Provided that the diameter of the action space  $\mathcal{A}$  is bounded, we can control the covering number  $\mathcal{N}(\mathcal{A}, \|\cdot\|_2, \varepsilon)$  and get a regret bound depending only on the dimension  $d$  and the time horizon  $T$ . It extends [2, Theorem 2] to general action spaces and improves on the constant factors inside the logarithm<sup>2</sup>.

**Corollary 2.** For  $d$ -linear bandit problems with rewards bounded in  $[-1, 1]$ , suppose that  $\mathcal{A} \subseteq \mathbf{B}_d(0, 1)$ , where  $\mathbf{B}_d(0, 1)$  is the  $d$ -dimensional closed Euclidean unit ball. Then, the TS cumulative regret is bounded as

$$\mathbb{E}[\text{Regret}(T)] \leq 2d \sqrt{T \log \left( \sqrt{2} + \frac{4\sqrt{T}}{d} \right)}.$$

*Proof.* The proof closely follows the technique from [2, Proof of Theorem 2]. Starting from Proposition 3, since  $\mathcal{A} \subseteq \mathbf{B}_d(0, 1)$  and  $\rho$  is the Euclidean distance, the logarithm of the covering number is bounded by  $d \log(1 + 2/\varepsilon)$  (see [16, Proof of lemma 5.13]). By setting  $\varepsilon = d/2\sqrt{T}$  and using properties of square roots and of logarithms, we obtain the claimed result.  $\square$

## VI. CONCLUSION

In this paper, we studied the Bayesian regret of the Thompson Sampling algorithm for bandit problems. Building on the analysis from Dong and Van Roy [2], we established bounds on the TS expected regret that hold even for problems with infinite and continuous action and parameter spaces. For the linear bandit problem, our analysis recovers the near-optimal rate of  $O(d\sqrt{T} \log T)$  from [2, Theorem 2] and improves the constants terms inside the logarithm. A natural direction for future work is to extend our results to other classes of bandit problems, such as the ‘‘semi-bandit’’ feedback problem.

<sup>2</sup>In [2, Theorem 2], the expected rewards are  $1/2$ -Lipschitz and the rewards in  $[-1/2, 1/2]$ . With the same assumptions, our bound reduces by a factor 2.

## REFERENCES

- [1] D. Russo and B. Van Roy, “An Information-Theoretic Analysis of Thompson Sampling,” Jun. 2015, number: arXiv:1403.5341 arXiv:1403.5341 [cs]. [Online]. Available: <http://arxiv.org/abs/1403.5341>
- [2] S. Dong and B. Van Roy, “An Information-Theoretic Analysis for Thompson Sampling with Many Actions,” May 2018, arXiv:1805.11845 [cs, math, stat]. [Online]. Available: <http://arxiv.org/abs/1805.11845>
- [3] D. Bouneffouf, I. Rish, and C. Aggarwal, “Survey on Applications of Multi-Armed and Contextual Bandits,” in *2020 IEEE Congress on Evolutionary Computation (CEC)*. Glasgow, United Kingdom: IEEE, Jul. 2020, pp. 1–8. [Online]. Available: <https://ieeexplore.ieee.org/document/9185782/>
- [4] N. Abe, A. W. Biermann, and P. M. Long, “Reinforcement Learning with Immediate Rewards and Linear Hypotheses,” *Algorithmica*, vol. 37, no. 4, pp. 263–293, Dec. 2003. [Online]. Available: <http://link.springer.com/10.1007/s00453-003-1038-1>
- [5] D. J. Russo, B. V. Roy, A. Kazerouni, I. Osband, and Z. Wen, “A Tutorial on Thompson Sampling,” *Foundations and Trends® in Machine Learning*, vol. 11, no. 1, pp. 1–96, Jul. 2018, publisher: Now Publishers, Inc. [Online]. Available: <https://www.nowpublishers.com/article/Details/MAL-070>
- [6] W. Chu, L. Li, L. Reyzin, and R. E. Schapire, “Contextual Bandits with Linear Payoff Functions,” *JMLR Workshop and Conference Proceedings*, 2011.
- [7] S. Agrawal and R. Jia, “Optimistic Posterior Sampling for Reinforcement Learning: Worst-Case Regret Bounds,” *Mathematics of Operations Research*, vol. 48, no. 1, pp. 363–392, Feb. 2023. [Online]. Available: <https://pubsonline.informs.org/doi/10.1287/moor.2022.1266>
- [8] D. J. Foster and A. Krishnamurthy, “Contextual bandits with surrogate losses: Margin bounds and efficient algorithms,” Nov. 2018, arXiv:1806.10745 [cs, stat]. [Online]. Available: <http://arxiv.org/abs/1806.10745>
- [9] —, “Efficient first-order contextual bandits: Prediction, allocation, and triangular discrimination.” *Advances in Neural Information Processing Systems*, 2021.
- [10] T. Zhang, “Feel-Good Thompson Sampling for Contextual Bandits and Reinforcement Learning,” Oct. 2021, arXiv:2110.00871 [cs, math, stat]. [Online]. Available: <http://arxiv.org/abs/2110.00871>
- [11] G. Neu, I. Olkhovskaia, M. Papini, and L. Schwartz, “Lifting the Information Ratio: An Information-Theoretic Analysis of Thompson Sampling for Contextual Bandits,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 9486–9498, Dec. 2022. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2022/hash/3d84d9b523e6e82916d496e58761002e-Abstract-Conference.html](https://proceedings.neurips.cc/paper_files/paper/2022/hash/3d84d9b523e6e82916d496e58761002e-Abstract-Conference.html)
- [12] W. R. Thompson, “On the likelihood that one unknown probability exceeds another in view of the evidence of two samples,” *Biometrika*, vol. 25, no. 3-4, pp. 285–294, Dec. 1933. [Online]. Available: <https://doi.org/10.1093/biomet/25.3-4.285>
- [13] D. Russo and B. Van Roy, “Learning to Optimize via Information-Directed Sampling,” Jul. 2017, arXiv:1403.5556 [cs]. [Online]. Available: <http://arxiv.org/abs/1403.5556>
- [14] O. Chapelle and L. Li, “An Empirical Evaluation of Thompson Sampling,” in *Advances in Neural Information Processing Systems*, vol. 24. Curran Associates, Inc., 2011. [Online]. Available: <https://proceedings.neurips.cc/paper/2011/hash/e53a0a2978c28872a4505bdb51db06dc-Abstract.html>
- [15] S. Dong, T. Ma, and B. V. Roy, “On the Performance of Thompson Sampling on Logistic Bandits,” *Conference on Learning Theory*, 2019.
- [16] R. van Handel, *Probability in High Dimension*. Princeton University, Dec. 2016, vol. APC 550 Lecture Notes.