
Generalized Thompson Sampling for Contextual Bandits

Lihong Li
Microsoft Research
Redmond, WA 98052
lihongli@microsoft.com

Abstract

Thompson Sampling, one of the oldest heuristics for solving multi-armed bandits, has recently been shown to demonstrate state-of-the-art performance. The empirical success has led to great interests in theoretical understanding of this heuristic. In this paper, we approach this problem in a way very different from existing efforts. In particular, motivated by the connection between Thompson Sampling and exponentiated updates, we propose a new family of algorithms called Generalized Thompson Sampling in the expert-learning framework, which includes Thompson Sampling as a special case. Similar to most expert-learning algorithms, Generalized Thompson Sampling uses a loss function to adjust the experts' weights. General regret bounds are derived, which are also instantiated to two important loss functions: square loss and logarithmic loss. In contrast to existing bounds, our results apply to quite general contextual bandits. More importantly, they quantify the effect of the "prior" distribution on the regret bounds.

1 Introduction

Thompson Sampling [18], one of the oldest heuristics for solving stochastic multi-armed bandits, embodies the principle of *probability matching*. Given a prior distribution over the underlying, unknown reward generating process as well as past observations of rewards, one can maintain a posterior distribution of which arm is optimal. Thompson Sampling then selects arms randomly according to the current posterior distribution.

While having been unpopular for decades, this algorithm was recently shown to be state-of-the-art in empirical studies, and has found success in important applications like news recommendation and online advertising [16, 10, 7, 14]. In addition, it has other advantages such as robustness to observation delay [7] and simplicity in implementation, compared to the dominant strategies based on upper confidence bounds (UCB).

Despite the empirical success, theoretical understanding of finite-time performance of Thompson Sampling has been limited until very recently. The first such result is provided by [2] for *non-contextual* K -armed bandits, who prove a nontrivial problem-dependent regret bound when the prior of an arm's expected reward is a Beta distribution. Later on, improved bounds are found for the same setting [11, 3], which match the asymptotic regret lower bound [12].

For contextual bandits [13], only two pieces of work are available, to the best of our knowledge. [4] analyze linear bandits, where a Gaussian prior is used on the weight vector space, and a Gaussian likelihood function is assumed for the reward function. The authors are able to show the regret grows on the order of $d\sqrt{T}$, which is only a \sqrt{d} factor away from a known matching lower bound [8]. In contrast, [15] establish an interesting connection between UCB-style analysis and the *Bayes risk* of Thompson Sampling, based on the probability-matching property. This observation allows the

authors to obtain Bayes risk bound based on a novel metric, known as *margin dimension*, of an arbitrary function class that essentially measures how fast upper confidence bounds decay.

All the existing work above relies critically either on advanced properties of the assumed prior distribution (such as in the case of Beta distributions), or on the assumption that the prior is correct (in the analysis of Bayes risk of [15]). Such analysis, although very interesting and important for better understanding Thompson Sampling, seems hard to be generalized to general (possibly nonlinear) contextual bandits. Furthermore, none of the existing theory is able to quantify the role of prior plays in controlling the regret, although in practice better domain knowledge is often available to construct good priors that should “accelerate” learning.

This paper attempts to address the limitations of prior work, from a very different angle. Based on a connection between Thompson Sampling and exponentiated update rules, we propose a family of contextual-bandit algorithms called Generalized Thompson Sampling in the expert-learning framework [6], where each expert corresponds to a contextual policy for arm selection. Similar to Thompson Sampling, Generalized Thompson Sampling is a randomized strategy, following an expert’s policy more often if the expert is more likely to be optimal. Different from Thompson Sampling, it uses a loss function to update the experts’ weights; Thompson Sampling is a special of Generalized Thompson Sampling when the logarithmic loss is used.¹

Regret bounds are then derived under certain conditions. The proof relies critically on a novel application of a “self-boundedness” property of loss functions in competitive analysis. The results are instantiated to the square and logarithmic losses, two important loss functions. Not only do these bounds apply to quite general sets of experts, but they also quantify the impact of the prior distribution on regret. These benefits come at a cost of a worse dependence on the number of steps. However, we believe it is possible to close the gap with a more involved analysis, and the connection between (Generalized) Thompson Sampling to expert-learning will likely lead to further interesting insights and algorithms in future work.

2 Preliminaries

Contextual bandits can be formulated as the following game between the learner and a stochastic environment. Let \mathcal{X} and \mathcal{A} be the sets of context and arms, and let $K = |\mathcal{A}|$. At step $t = 1, 2, \dots, T$:

- Learner observes the context $x_t \in \mathcal{X}$, where x_t can be chosen by an adversary.
- Learner selects arm $a_t \in \mathcal{A}$, and receives reward $r_t \in \{0, 1\}$, with expectation $\mu(x_t, a_t)$.

Note that the setup above allows the contexts to be chosen by an adversary, which is a more general setting than typical contextual bandits [13]. The reader may notice we require the reward to be binary, instead of being in $[0, 1]$. This choice will make our exposition simpler, without sacrificing loss of generality. Indeed, as also suggested by [2], if reward $r \in (0, 1)$ is received, one can convert it into a binary pseudo-reward $\tilde{r} \in \{0, 1\}$ as follows: let \tilde{r} be 1 with probability r , and 0 otherwise. Clearly, the bandit process remains essentially the same, with the same optimal expert and regrets.

Motivated by prior work on Thompson Sampling with parametric function classes [7], we allow the learner to have access to a set of experts, $\mathcal{E} = \{\mathcal{E}_1, \dots, \mathcal{E}_N\}$, each one of them makes predicts about the average reward $\mu(x, a)$. Let f_i be the associated prediction function of expert \mathcal{E}_i . Its arm-selection policy in context x is simply the greedy policy with respect to the reward predictions: $\mathcal{E}_i(x) = \max_{a \in \mathcal{A}} f_i(x, a)$. This setting can naturally be used to capture the use of parametric function classes: for example, when generalized linear models are used to predict $\mu(x, a)$ [10, 7], each weight vector is an expert. The only difference is that our framework works with a discrete set of experts. Using a covering device, however, it is possible to approximate a continuous function class by a finite set of cardinality N , where N is the covering number.

We define the T -step average regret of the learner by

$$R(T) = \max_{1 \leq i \leq N} \sum_{t=1}^T \mu(x_t, \mathcal{E}_i(x_t)) - \mathbf{E} \left[\sum_{t=1}^T \mu(x_t, a_t) \right], \quad (1)$$

¹It should be emphasized that, in this paper, we use the loss function to measure how well an expert *predicts* the average reward, given the context and the selected arm. In general, the loss function and the reward may be completely unrelated. Details are given later.

where the expectation refers to the possible randomization of the learner in selecting a_t . As in all existing analysis for Thompson Sampling, we make the realization assumption that one of the experts, $\mathcal{E}^* \in \mathcal{E}$, correctly predicts the average reward. Without loss of generality, let \mathcal{E}_1 be this expert; in other words, $f_1(x, a) \equiv \mu(x, a)$. Clearly, \mathcal{E}_1 is the reward-maximizing expert, so $R(T) = \sum_t \mu(x_t, \mathcal{E}_1(x_t)) - \mathbf{E}[\sum_t \mu(x_t, a_t)]$.

With the notation above, Thompson Sampling can be described as follows. It requires as input a ‘‘prior’’ distribution $\mathbf{p} = (p_1, \dots, p_N) \in \mathbb{R}_+^N$ over the experts, where $\|\mathbf{p}\|_1 = 1$. Intuitively, p_i may be interpreted as the prior probability that \mathcal{E}_i is the reward-maximizing expert. The algorithm starts with the first ‘‘posterior’’ distribution $\mathbf{w}_1 = (w_{1,1}, \dots, w_{N,1})$ where $w_{i,1} = p_i$. At step t , the algorithm samples an expert based on the posterior distribution \mathbf{w}_t and follows that expert’s policy to choose action. Upon receiving the reward, the weights are updated by $w_{i,t+1} \propto w_{i,t} \exp(-\ell(f_i(x_t, a_t), r_t))$, where $\ell(f, r)$ is the negative log-likelihood.

Finally, one can assume the optimal expert, \mathcal{E}^* , is drawn from an *unknown* prior distribution, $\mathbf{p}^* = (p_1^*, \dots, p_N^*)$. The expected T -step *Bayes regret* can then be defined: $R(T, \mathbf{p}^*) \stackrel{\text{def}}{=} \mathbf{E}_{\mathcal{E}^* \sim \mathbf{p}^*} [R(T)]$. It should be noted that the Bayes risk considered by other authors [15] is just $R(T, \mathbf{p})$, where \mathbf{p} is the prior used by Thompson Sampling. In general, the true prior \mathbf{p}^* is unknown, so $\mathbf{p} \neq \mathbf{p}^*$. We believe the Bayes risk defined with respect to \mathbf{p}^* is more reasonable in light of the almost inevitable misspecification of priors in practice.

3 Generalized Thompson Sampling

An observation with Thompson Sampling from the previous section is that its Bayes update rule can be viewed as an exponentiated update with the logarithmic loss (see also [6]). After receiving a reward, each expert is penalized for the mismatch in its prediction (f_i) and the observed reward, and the penalty happens to be the logarithmic loss in Thompson Sampling. Therefore, in principle, one can use other loss function to get a more general family of algorithms. In fact, *none* of the existing regret analyses [2, 3, 4, 11] relies on the interpretation that \mathbf{w}_t are meant to be Bayesian posteriors, and yet manages to show strong regret bound for Thompson Sampling.² The above observations suggest the promising performance of Thompson Sampling is not due to its Bayesian nature, and also motivates us to develop a more general family of algorithms known as *Generalized Thompson Sampling*.

We denote by $\ell(\hat{r}, r)$ the loss incurred by reward prediction \hat{r} when the observed reward is r . Generalized Thompson Sampling performs exponentiated updates to adjust experts’ weights, and follows a randomly selected expert when making decisions, similar to Thompson Sampling. In addition, the algorithm also allows mixing of the exponentially weighted distribution and a uniform distribution controlled by γ . The pseudocode is given in Algorithm 1.

Algorithm 1 Generalized Thompson Sampling

Input: $\eta > 0, \gamma > 0, \{\mathcal{E}_1, \dots, \mathcal{E}_N\}$, and prior \mathbf{p}

Initialize posterior: $\mathbf{w}_1 \leftarrow \mathbf{p}; W_1 \leftarrow \|\mathbf{w}_1\|_1 = 1$

for $t = 1, \dots, T$ **do**

 Receive context $x_t \in \mathcal{X}$

 Select arm a_t according to the mixture probabilities: for each a

$$\Pr(a) = (1 - \gamma) \sum_{i=1}^N \frac{w_{i,t} \mathbb{I}(\mathcal{E}_i(x_t) = a)}{W_t} + \frac{\gamma}{K}$$

 Observe reward r_t , and updates weights:

$$\forall i : w_{i,t+1} \leftarrow w_{i,t} \cdot \exp(-\eta \cdot \ell(f_i(x_t, a_t), r_t)); \quad W_{t+1} \leftarrow \|\mathbf{w}_{t+1}\|_1 = \sum_i w_{i,t+1}$$

end for

²The analysis of [15] is different since the metric (Bayes risk) is defined with respect to the prior.

Clearly, Generalized Thompson Sampling includes Thompson Sampling as a special case, by setting $\eta = 1$, $\gamma = 0$, and ℓ to be the logarithmic loss: $\ell(\hat{r}, r) = \mathbb{I}(r = 1) \ln 1/\hat{r} + \mathbb{I}(r = 0) \ln 1/(1 - \hat{r})$. Another loss function considered in this paper is the square loss: $\ell(\hat{r}, r) = (\hat{r} - r)^2$.

4 Analysis

For convenience, the analysis here uses the following shorthand notation:

- The history of the learner up to step t is $F_t \stackrel{\text{def}}{=} (x_1, a_1, r_1, \dots, x_{t-1}, a_{t-1}, r_{t-1}, x_t)$.
- The *immediate regret* of expert \mathcal{E}_i in context x is $\Delta_i(x) \stackrel{\text{def}}{=} \mu(x, \mathcal{E}_i(x)) - \mu(x, \mathcal{E}_1(x))$.
- The *normalized weight* at step t is $\bar{w}_t \stackrel{\text{def}}{=} \mathbf{w}_t / W_t$.
- The *shifted loss* incurred by expert \mathcal{E}_i in triple (x, a, r) is denoted by $\hat{l}_i(r|x, a) \stackrel{\text{def}}{=} \ell(f_i(x, a), r) - \ell(f_1(x, a), r)$. In particular, define $\hat{l}_{i,t} \stackrel{\text{def}}{=} \hat{l}_i(r_t|x_t, a_t)$. In other words, \hat{l}_i is the loss relative to the best expert (\mathcal{E}_1), and *can be negative*.
- The *average shifted loss* at step t is $\bar{l}_t \stackrel{\text{def}}{=} \mathbf{E}_{r_t, a_t | F_t} \left[\sum_i \bar{w}_{i,t} \hat{l}_i(r_t|x_t, a_t) \right]$.

4.1 Main Theorem

Clearly, conditions are needed to relate the loss function to the regret. Our results need the following assumptions:

- (C1) (Consistency) For all $(x, a) \in \mathcal{X} \times \mathcal{A}$, $\mathbf{E}_{r|x,a} \left[\hat{l}_i(r|x, a) \right] \geq 0$.
- (C2) (Informativeness) There exists a constant $\kappa_1 \in \mathbb{R}_+$ such that $\Delta_i(x_t) \leq \kappa_1 \sqrt{\bar{l}_t}$.
- (C3) (Boundedness) The shifted loss \hat{l}_i assumes values in $[-1, 1]$.
- (C4) (Self-boundedness) There exists a constant $\kappa_2 \in \mathbb{R}_+$ such that, for all $(x, a) \in \mathcal{X} \times \mathcal{A}$, $\mathbf{E}_{r|x,a} \left[\hat{l}_i(r|x, a)^2 \right] \leq \kappa_2 \mathbf{E}_{r|x,a} \left[\hat{l}_i(r|x, a) \right]$; namely, the second moment is bounded, up to a constant, by the first moment of the shifted loss.

Theorem 1 *Under Conditions C1 and C2, the expected T -step regret of Generalized Thompson Sampling is*

$$R(T) \leq \kappa_1 \sqrt{T \cdot \mathbf{E} \left[\sum_{t=1}^T \bar{l}_t \right]} + \gamma T.$$

Proof. The expected T -step regret may be rewritten more explicitly, and then bounded, as follows:

$$\begin{aligned} R(T) &\leq \mathbf{E} \left[\sum_{t=1}^T \left((1 - \gamma) \sum_{i=1}^N \bar{w}_{i,t} \Delta_i(x_t) + \gamma \right) \right] \\ &= (1 - \gamma) \mathbf{E} \left[\sum_t \sum_i \bar{w}_{i,t} \Delta_i(x_t) \right] + \gamma T \\ &\leq (1 - \gamma) \mathbf{E} \left[\sum_t \sum_i \bar{w}_{i,t} \kappa_1 \sqrt{\mathbf{E}_{r_t, a_t | F_t} \left[\hat{l}_i(r_t|x_t, a_t) \right]} \right] + \gamma T \\ &= \kappa_1 (1 - \gamma) \mathbf{E} \left[\sum_t \sum_i \bar{w}_{i,t} \sqrt{\mathbf{E}_{r_t, a_t | F_t} \left[\hat{l}_i(r_t|x_t, a_t) \right]} \right] + \gamma T \\ &\leq \kappa_1 (1 - \gamma) \mathbf{E} \left[\sum_t \sqrt{\sum_i \bar{w}_{i,t} \mathbf{E}_{r_t, a_t | F_t} \left[\hat{l}_i(r_t|x_t, a_t) \right]} \right] + \gamma T \\ &\leq \kappa_1 (1 - \gamma) \sqrt{T \cdot \mathbf{E} \left[\sum_t \bar{l}_t \right]} + \gamma T. \end{aligned}$$

□

Now the question becomes one of bounding the expected total shifted loss, $\mathbf{E} [\sum_t \bar{l}_t]$. This problem is tackled by the following key lemma, which makes use of the self-boundedness property of the loss function. The lemma may be of interest on its own. Similar properties were used in [1] in a very different way.

Lemma 1 *Under Conditions C3 and C4, with η chosen to be $(2(e-2)\kappa_2)^{-1}$, the expected total shifted loss of Generalized Thompson Sampling is bounded by a constant independent of T :*

$$\sum_{t=1}^T \bar{l}_t \leq 4(e-2)\kappa_2 \ln \frac{1}{p_1}.$$

Proof. First, observe that if the shifted loss $\hat{l}_{i,t}$ is used in Generalized Thompson Sampling to replace the loss $\ell(f_i(x_t, a_t), r_t)$, the algorithm behaves identically. The rest of the proof uses this fact, pretending Generalized Thompson Sampling uses $\hat{l}_{i,t}$ for weight updates.

For any step t , the weight sum changes according to

$$\begin{aligned} \ln \frac{W_{t+1}}{W_t} &= \ln \left(\sum_i \bar{w}_{i,t} \exp(-\eta \hat{l}_{i,t}) \right) \\ &\leq \ln \left(\sum_i \bar{w}_{i,t} \left(1 - \eta \hat{l}_{i,t} + (e-2)\eta^2 \hat{l}_{i,t}^2 \right) \right) \\ &= \ln \left(1 - \eta \sum_i \bar{w}_{i,t} \hat{l}_{i,t} + (e-2)\eta^2 \sum_i \bar{w}_{i,t} \hat{l}_{i,t}^2 \right) \\ &\leq -\eta \sum_i \bar{w}_{i,t} \hat{l}_{i,t} + (e-2)\eta^2 \sum_i \bar{w}_{i,t} \hat{l}_{i,t}^2, \end{aligned}$$

where the first inequality is due to Condition C3 and the inequality $e^x \leq 1 - x + (e-2)x^2$ for $x \in [-1, 1]$; the second inequality is due to the inequality $\ln(1-x) < x$ for $x < 1$.

Conditioned on the observed context and selected arm at step t , we take expectation of the above expressions, with respect to the randomization in observed reward, leading to

$$\mathbf{E}_{r|F_t, a_t} \left[\ln \frac{W_{t+1}}{W_t} \right] \leq -\eta \sum_i \bar{w}_{i,t} \mathbf{E}_{r|F_t, a_t} [\hat{l}_{i,t}] + (e-2)\eta^2 \sum_i \bar{w}_{i,t} \mathbf{E}_{r|F_t, a_t} [\hat{l}_{i,t}^2].$$

Condition C4 then implies

$$\mathbf{E}_{r|F_t, a_t} \left[\ln \frac{W_{t+1}}{W_t} \right] \leq -\eta (1 - (e-2)\eta\kappa_2) \sum_i \bar{w}_{i,t} \mathbf{E}_{r|F_t, a_t} [\hat{l}_{i,t}].$$

Setting $\eta = (2(e-2)\kappa_2)^{-1}$ gives

$$\mathbf{E}_{r|F_t, a_t} \left[\ln \frac{W_{t+1}}{W_t} \right] \leq -\frac{1}{4(e-2)\kappa_2} \sum_i \bar{w}_{i,t} \mathbf{E}_{r|F_t, a_t} [\hat{l}_{i,t}].$$

The above inequality holds for any a_t , so also holds in expectation if a_t is randomized:

$$\mathbf{E}_{a_t, r_t|F_t} \left[\ln \frac{W_{t+1}}{W_t} \right] \leq -\frac{1}{4(e-2)\kappa_2} \sum_i \bar{w}_{i,t} \mathbf{E}_{a_t, r_t|F_t} [\hat{l}_{i,t}].$$

Finally, summing the left-hand side over $t = 1, 2, \dots, T$ gives

$$\mathbf{E} \left[\frac{W_{T+1}}{W_1} \right] \leq -\frac{1}{4(e-2)\kappa_2} \sum_t \sum_i \bar{w}_{i,t} \mathbf{E}_{a_t, r_t|F_t} [\hat{l}_{i,t}] = -\frac{1}{4(e-2)\kappa_2} \mathbf{E} \left[\sum_t \bar{l}_t \right],$$

which implies

$$\mathbf{E} \left[\sum_t \bar{l}_t \right] \leq 4(e-2)\kappa_2 \mathbf{E} \left[\frac{W_1}{W_{T+1}} \right] \leq 4(e-2)\kappa_2 \ln \frac{1}{w_{1,1}} = 4(e-2)\kappa_2 \ln \frac{1}{p_1}.$$

The last inequality above follows from the observation that $\hat{l}_{1,t} \equiv 0$, and that $w_{1,t} \equiv w_{1,1}$. \square

The following corollary follows directly from Theorem 1 and Lemma 1:

Corollary 1 *Under the conditions of Theorem 1 and Lemma 1, the expected T -step regret of Generalized Thompson Samplings is at most*

$$\sqrt{4\kappa_2(e-2)\kappa_1(1-\gamma)} \sqrt{T \cdot \ln \frac{1}{p_1}} + \gamma T.$$

The next corollary considers the Bayes regret, $R(T, \mathbf{p}^*)$, with an unknown, true prior \mathbf{p}^* :

Corollary 2 *If the optimal expert is sampled from distribution \mathbf{p}^* , the Bayes regret is at most*

$$R(T, \mathbf{p}^*) \leq \sqrt{4\kappa_2(e-2)\kappa_1(1-\gamma)} \sqrt{T (H(\mathbf{p}^*) + d_{\text{KL}}(\mathbf{p}^*, \mathbf{p}))} + \gamma T,$$

where $H(\mathbf{p}^*)$ and $d_{\text{KL}}(\mathbf{p}^*, \mathbf{p})$ are the standard entropy and KL-divergence.

Proof. We have

$$\begin{aligned} R(T, \mathbf{p}^*) &\leq \sum_{i=1}^N p_i^* \left(\sqrt{4\kappa_2(e-2)\kappa_1(1-\gamma)} \sqrt{T \cdot \ln \frac{1}{w_{i,1}}} + \gamma T \right) \\ &= \sqrt{4\kappa_2(e-2)\kappa_1(1-\gamma)} \sqrt{T} \sum_i \left(w_i^* \sqrt{\ln \frac{1}{w_i}} \right) + \gamma T \\ &\leq \sqrt{4\kappa_2(e-2)\kappa_1(1-\gamma)} \sqrt{T} \sqrt{\sum_i \left(w_i^* \ln \frac{1}{w_i} \right)} + \gamma T \\ &= \sqrt{4\kappa_2(e-2)\kappa_1(1-\gamma)} \sqrt{T (H(\mathbf{p}^*) + d_{\text{KL}}(\mathbf{p}^*, \mathbf{p}))} + \gamma T, \end{aligned}$$

where the inequalities are due to Corollary 1 and Jensen's inequality, respectively. \square

4.2 Square Loss

We start with the simpler case of square loss. It clearly satisfies Condition C3. Condition C1 holds because of the following well-known fact:

$$\mathbf{E}_{r|x,a} \left[\hat{l}_i(r|x,a) \right] = \mathbf{E}_{r|x,a} \left[(f_i(x,a) - r)^2 - (f_1(x,a) - r)^2 \right] = (f_i(x,a) - f_1(x,a))^2 \geq 0.$$

Conditions C2 and C4 are also satisfied with $\kappa_1 = \sqrt{\frac{2K}{\gamma}}$ and $\kappa_2 = 4$, from prior work [1]. Plugging these values in Corollary 1 and choosing $\gamma = \Theta \left(\sqrt[3]{K/T} \right)$, we obtain the regret bound of $O \left(\sqrt{\ln \frac{1}{p_1}} K^{1/3} T^{2/3} \right)$, and the Bayes regret bound of $O \left(\sqrt{H(\mathbf{p}^* + d_{\text{KL}}(\mathbf{p}^*, \mathbf{p}))} K^{1/3} T^{2/3} \right)$.

4.3 Logarithmic Loss

For logarithmic loss, we assume the shifted loss of all experts are bounded in $[-\beta/2, \beta/2]$ for some constant $\beta \in \mathbb{R}_+$, so that one can normalize the shifted logarithmic loss to the range of $[-1, 1]$ by defining:

$$l_i(r|x,a) = \frac{\mathbb{I}(r=1)}{\beta} \ln \frac{1}{f_i(x,a)} + \frac{\mathbb{I}(r=0)}{\beta} \ln \frac{1}{1-f_i(x,a)}. \quad (2)$$

This assumption can usually be satisfied in practice, and seems necessary to derive finite-time guarantees. Note that this assumption is slightly weaker than the more common assumption that the logarithmic loss itself is bounded (e.g., [9]).

We now verify all necessary conditions. Condition C1 follows from the well-known fact that the expectation of logarithmic loss between the true expert and another is their KL-divergence,

$$\mathbf{E}_{r|x,a} [\hat{l}_i(r|x,a)] = \frac{1}{\beta} d_{\text{KL}}(f_1(x,a), f_i(x,a)) \quad (3)$$

which is in turn non-negative.

Condition C2 is verified in the following lemma:

Lemma 2 *For the loss function defined in Equation (2), one has*

$$\Delta_i(x) \leq K \sqrt{\frac{2\beta}{\gamma}} \sqrt{\mathbf{E}_{r,a|x} [\hat{l}_i(r|x,a)]}.$$

Proof. We have the following:

$$\begin{aligned} \Delta_i(x) &= f_1(x, \mathcal{E}_1(x)) - f_1(x, \mathcal{E}_i(x)) \\ &\leq |f_1(x, \mathcal{E}_1(x)) - f_i(x, \mathcal{E}_1(x))| + |f_1(x, \mathcal{E}_i(x)) - f_i(x, \mathcal{E}_i(x))| \\ &\leq \sqrt{2d_{\text{KL}}(f_1(x, \mathcal{E}_1(x)), f_i(x, \mathcal{E}_1(x)))} + \sqrt{2d_{\text{KL}}(f_1(x, \mathcal{E}_i(x)), f_i(x, \mathcal{E}_i(x)))} \\ &\leq \sum_a \sqrt{2d_{\text{KL}}(f_1(x,a), f_i(x,a))} \\ &\leq \sqrt{2K} \sqrt{\sum_a d_{\text{KL}}(f_1(x,a), f_i(x,a))} \\ &\leq \sqrt{\frac{2K}{\gamma/K}} \sqrt{\mathbf{E}_{a|x} d_{\text{KL}}(f_1(x,a), f_i(x,a))} \\ &= K \sqrt{\frac{2\beta}{\gamma}} \sqrt{\mathbf{E}_{r,a|x} [\hat{l}_i(r|x,a)]}, \end{aligned}$$

where the first inequality is due to the triangle inequality; the second inequality is due to Pinsker's inequality; the fourth inequality is due to Jensen's inequality; the fifth inequality is from the fact that each arm is selected with probability at least γ/K ; the last equality is from Equation (3). \square

Condition C3 is immediately satisfied by the normalization of $1/\beta$ in the definition of l_i above.

Condition C4 is the most difficult one to verify. To the best of our knowledge, such a result for logarithmic loss is not found in literature and can be of independent interest. For example, it implies that the analysis of [1] for square loss also applies to the logarithmic loss. The following lemma states the result more formally. Its proof, which is rather technical, is left to the appendix.

Lemma 3 *For the loss function defined in Equation (2), there exists some constant $\kappa_2 = O(1)$ such that Condition C4 holds.*

With all four conditions verified, we can apply results in Section 4.1 to reach the regret bound of $O\left(K^{\frac{2}{3}} \beta^{\frac{1}{3}} T^{\frac{2}{3}} \sqrt{\ln \frac{1}{p_1}}\right)$, and the Bayes regret bound of $O\left(K^{\frac{2}{3}} \beta^{\frac{1}{3}} T^{\frac{2}{3}} \sqrt{H(w^*) + d_{\text{KL}}(w^*, w)}\right)$.

5 Discussions

In this paper, we propose a new family of algorithms, Generalized Thompson Sampling, and analyze its regret in the expert-learning framework. Our regret analysis provides a promising alternative to understanding the strong performance of Thompson Sampling, an interesting and pressing research problem raised by its recent empirical success. Compared to existing analysis in the literature, it has the following benefits. First, the results apply more generally to a set of experts, rather than making specific modeling assumptions about the prior and likelihood. Second, the analysis quantifies how the (not necessarily correct prior \mathbf{p}) affects the regret bound, as well as the Bayes regret when optimal experts are drawn from an unknown prior \mathbf{p}^* . Similar to PAC-Bayes bounds, these results combine the benefits of good priors and the robustness of frequentist approaches.

Our proof for Generalized Thompson Sampling is inspired by the online-learning literature [6]. However, a new technique is needed to prove the critical Lemma 1, which relies on self-boundedness of a loss function. A similar property is shown by [1] for square loss only, and is used in a very different way. The self-boundedness of logarithmic loss (Lemma 3) appears new, to the best of our knowledge, and may be of independent interest.

Generalized Thompson Sampling bears some similarities to the Regressor Elimination (RE) algorithm [1]. A crucial difference is that RE requires a computationally expensive operation of computing a “balanced” distribution over experts, in order to control variance in the elimination process. In contrast, our algorithm is computationally much cheaper. The operations of Generalized Thompson Sampling are also related to EXP4 [5], which uses unbiased, importance-weighted reward estimates to do exponentiated updates of expert weights. In practice, it seems more natural to use prediction loss of an expert to adjust its weight, rather than using the reward signals directly [10, 7].

While we have focused on the case of finitely many experts, the setting is motivated by the more realistic case when the set \mathcal{E} of experts is continuous [10, 7, 4]. The discrete case considered here may be thought of as an approximation to the continuous case, using a covering device. We expect similar results to hold with N replaced by the covering number of the class.

This work suggests a few interesting directions for future work. The first is to close the gap between the current $O(T^{2/3})$ bound and the best problem-independent bound $O(\sqrt{T})$ for contextual bandits. The second is to extend the analysis here to continuous expert classes, and more importantly to the agnostic (non-realizable) case. Finally, it is interesting to use the regret analysis of (Generalized) Thompson Sampling to obtain performance guarantees for its reinforcement-learning analogues (*e.g.*, [17]).

A Proof of Lemma 3: Self-boundedness of Logarithmic Loss

This section proves Lemma 3, regarding self-boundedness of logarithmic loss, in the sense described in Condition C4. The analysis here does not involve step t and the corresponding context and selected arm. We therefore simplify notation as follows: the true expert \mathcal{E}_1 predicts $f \in (0, 1)$, and the other expert \mathcal{E}_i predicts $g \in (0, 1)$. The binary reward is then a Bernoulli random variable with success rate f . The shifted logarithmic loss of \mathcal{E}_2 is given by

$$\hat{l} = \frac{\mathbb{I}(r=1)}{\beta} \ln \frac{f}{g} + \frac{\mathbb{I}(r=0)}{\beta} \ln \frac{1-f}{1-g}.$$

The first two moments of the random variable \hat{l} are given by:

$$\begin{aligned} M_1 &\stackrel{\text{def}}{=} \mathbf{E}_r [\hat{l}] = \frac{d_{\text{KL}}(f, g)}{\beta} = \frac{f}{\beta} \ln \frac{f}{g} + \frac{1-f}{\beta} \ln \frac{1-f}{1-g} \\ M_2 &\stackrel{\text{def}}{=} \mathbf{E}_r [\hat{l}^2] = \frac{f}{\beta^2} \left(\ln \frac{f}{g} \right)^2 + \frac{1-f}{\beta^2} \left(\ln \frac{1-f}{1-g} \right)^2. \end{aligned}$$

Define

$$F(g) \stackrel{\text{def}}{=} \frac{M_2 - M_1^2}{M_1},$$

the ratio between variance and expectation of \hat{l} , as a function of g . Our goal is to show that $F(g)$ is bounded by a constant, independent of f and g . It will then follow that M_2/M_1 is also bounded by a constant since $M_2/M_1 = F(g) + M_1 \leq F(g) + 1$.

Taking the derivative of F , one obtains

$$F'(g) = -C(f, g) \ln \frac{f(1-g)}{g(1-f)} \left((f+g) \ln \frac{f}{g} + (2-f-g) \ln \frac{1-f}{1-g} \right)$$

for some function $C(f, g) > 0$. It can be verified, by rather tedious calculations, that there exists some $g_0 \in (0, 1)$ such that $F'(g) \leq 0$ for $g < g_0$ and $F'(g) \geq 0$ for $g > g_0$. So $F(g)$ is maximized by making g close to either 0 or 1. It then follows, again by rather tedious calculations, that $F(g) = O(1)$, using the assumption that the log-ratios (that is, shifted loss) are bounded by $[-\beta/2, \beta/2]$.

References

- [1] Alekh Agarwal, Miroslav Dudík, Satyen Kale, John Langford, and Robert E. Schapire. Contextual bandit learning under the realizability assumption. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics (AISTATS-12)*, 2012.
- [2] Shipra Agrawal and Navin Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proceedings of the Twenty-Fifth Annual Conference on Learning Theory (COLT-12)*, pages 39.1–39.26, 2012.
- [3] Shipra Agrawal and Navin Goyal. Further optimal regret bounds for Thompson sampling. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS-13)*, 2013.
- [4] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of Thirtieth International Conference on Machine Learning (ICML-13)*, 2013.
- [5] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- [6] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [7] Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems 24 (NIPS-11)*, pages 2249–2257, 2012.
- [8] Wei Chu, Lihong Li, Lev Reyzin, and Robert E. Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS-11)*, pages 208–214, 2011.
- [9] Sarah Filippi, Olivier Cappe, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems 23 (NIPS-10)*, pages 586–594, 2011.
- [10] Thore Graepel, Joaquin Quinonero Candela, Thomas Borchert, and Ralf Herbrich. Web-scale Bayesian click-through rate prediction for sponsored search advertising in Microsoft’s Bing search engine. In *Proceedings of the Twenty-Seventh International Conference on Machine Learning (ICML-10)*, pages 13–20, 2010.
- [11] Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *Proceedings of the Twenty-Third International Conference on Algorithmic Learning Theory (ALT-12)*, pages 199–213, 2012.
- [12] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- [13] John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. In *Advances in Neural Information Processing Systems 20*, pages 1096–1103, 2008.
- [14] Benedict C. May, Nathan Korda, Anthony Lee, and David S. Leslie. Optimistic Bayesian sampling in contextual-bandit problems. *Journal of Machine Learning Research*, 13:2069–2106, 2012.
- [15] Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 2013. accepted.
- [16] Steven L. Scott. A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26:639–658, 2010.
- [17] Malcolm J. A. Strens. A Bayesian framework for reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML-00)*, pages 943–950, 2000.
- [18] 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.