

---

# Distilled Thompson Sampling: Practical and Efficient Thompson Sampling via Imitation Learning

---

**Hongseok Namkoong\***  
Columbia Business School  
namkoong@gsb.columbia.edu

**Samuel Daulton\***  
Facebook  
sdaulton@fb.com

**Eytan Bakshy**  
Facebook  
ebakshy@fb.com

## Abstract

Thompson sampling (TS) has emerged as a robust technique for contextual bandit problems. However, TS requires posterior inference and optimization for action generation, prohibiting its use in many internet applications where latency and ease of deployment are of concern. We propose a novel imitation-learning-based algorithm that distills a TS policy into an explicit policy representation by performing posterior inference and optimization offline. The explicit policy representation enables fast online decision-making and easy deployment in mobile and server-based environments. Our algorithm iteratively performs offline batch updates to the TS policy and learns a new imitation policy. Since we update the TS policy with observations collected under the imitation policy, our algorithm emulates an off-policy version of TS. Our imitation algorithm guarantees Bayes regret comparable to TS, up to the sum of single-step imitation errors. We show these imitation errors can be made arbitrarily small when unlabeled contexts are cheaply available, which is the case for most large-scale internet applications. Empirically, we show that our imitation policy achieves comparable regret to TS, while reducing decision-time latency by over an order of magnitude.

## 1 Introduction

In the past decade, Thompson sampling [54] has emerged as a powerful algorithm for contextual bandit problems. The underlying principle is simple: an action is chosen with probability proportional to it being optimal under the current posterior distribution. Driven by the algorithm’s strong empirical performance [50, 15, 39] and rigorous performance guarantees [30, 7, 8, 25, 28, 47, 3], Thompson sampling has gained popularity in a broad range of applications including revenue management [21], internet advertising [26, 5, 49], and recommender systems [31].

Despite its conceptual simplicity and strong performance, Thompson sampling can be difficult to deploy. Concretely, choosing an action with Thompson sampling consists of two steps: *posterior sampling* and *optimization*. *Posterior sampling* requires evaluating a potentially large number of actions from a well-calibrated probabilistic model. Uncertainty calibration is crucial for empirical performance [45], since a well-calibrated model allows optimally trading off exploration and exploitation. In this regard, large-scale (probabilistic) machine learning models based on deep networks show much promise as they can adaptively learn good feature representations. However, sampling from these probabilistic models can be demanding in terms of both computation and memory. While approximate inference methods with better runtime characteristics exist, they tend to produce poorly calibrated uncertainty estimates that translate into poorer empirical performance on bandit tasks [45]. The second step, *optimization*, solves for a reward-optimizing action under the posterior sample. This can be prohibitively computationally expensive when the action space is large or continuous. For

---

\*Equal contribution

example, an advertising platform deciding to match advertisers to users has to solve combinatorial optimization problems in real-time in order to run Thompson sampling [38].

These challenges are pronounced in mobile applications; as of 2018, an estimated 52.2% of web traffic was generated by mobile devices [52]. Mobile applications require decisions to be made in a fast and memory-efficient manner, and on-device decision-making is critical to good user experience in domains such as adaptive video streaming [37] and social media ranking [42]. However, the majority of internet-connected mobile devices have limited memory, and rely on low-end processors that are orders of magnitude slower than server-grade devices [10, 64]. Another practical consideration is the technical debt of implementing contextual bandit algorithms in large-scale internet services. The online nature of the complex routines required by Thompson sampling leads the overall system to be cumbersome and hard to debug, making reliable software development challenging.

Motivated by challenges in implementing and deploying Thompson sampling in real production systems, we develop and analyze a method that maintains an explicit policy representation designed to imitate Thompson sampling, which allows efficient action generation without real-time posterior inference or numerical optimization. In order to avoid complex and computationally demanding routines *online*, our algorithm simulates a Thompson sampling policy *offline* and learns an explicit policy parameterization that mimics the behavior of Thompson sampling. As in many production systems, we consider the setting where the policy is updated offline with batches of new data [27]. At each period, the imitation learning problem can be efficiently solved using stochastic gradient-based methods, as it is equivalent to a log likelihood maximization (MLE) problem where the goal is to find a policy parameterization maximizing the likelihood of observing the actions generated by Thompson sampling. With this distilled policy in hand, actions can be generated efficiently by sampling from this parameterized distribution, conditional on the observed context; if we use neural networks to parameterize our imitation policy, this corresponds to a single forward pass. By virtue of moving complex numerical routines offline, our imitation procedure allows arbitrary, state-of-the-art Bayesian models and optimization solvers to be used even in latency and memory sensitive environments, and can be easily implemented using modern machine learning pipelines [23, 22].

Since our updates to the Thompson sampling policy are based on observations generated by the imitation policy, our algorithm emulates an *off-policy* version of Thompson sampling. Our main theoretical result (Section 4) establishes that our imitation algorithm maintains performance comparable to the true on-policy Thompson sampling policy, up to the sum of single-step imitation errors. In particular, our results preclude the possibility of small deviations between our imitation policy and Thompson sampling magnifying over time. We prove that each single-period imitation error term can be controlled—with a sufficiently rich imitation model—at the rate  $O_p(1/\sqrt{N})$ , where  $N$  is the number of (potentially unlabeled, meaning those without corresponding actions or rewards) available contexts. Combining this with our regret bound, our imitation algorithm achieves performance comparable to Thompson sampling up to  $O_p(T/\sqrt{N})$ -error. Despite the seemingly linear regret, we often have  $T \ll \sqrt{N}$  in internet applications where databases with entity features provide abundant unlabeled contexts, often in the order of hundreds of millions. In contrast, the number of model updates (horizon  $T$ ) is relatively small, in tens or hundreds, due to complexities of reliable software deployment and nonstationary user behavior. In practical problem settings where contexts are abundant, our results show that the imitation policy enjoys Bayes regret comparable to that of Thompson sampling, achieving optimal (gap-independent) regret in the wealth of examples where Thompson sampling is known to be optimal.

We empirically evaluate our imitation algorithm on several benchmark problems and a real-world dataset for selecting optimal video transcoding configurations (Section 5). In all of our experiments, we find that that our imitation algorithm performs as well as Thompson sampling in terms of cumulative regret, while reducing decision-time latency by an order of magnitude.

**Related work** Many authors have showed regret bounds for TS [6, 7, 30, 25, 28, 47, 3, 48]; we refer the reader to the recent tutorial [48] for a comprehensive overview. We build on the insights of Russo and Van Roy [47], and show that our imitation algorithm retains the advantageous properties of TS, achieving (gap-independent) Bayes regret comparable to the *best* UCB algorithm.

The performance of explore-exploit algorithms like TS depend on having access to well-calibrated probabilistic predictions. Obtaining a balance between predictive accuracy, time, and space can be challenging in the context of large datasets with over-parameterized models. Exact posterior sampling from even the simplest Gaussian linear models has a time complexity of  $O(n^2)$ , where  $n$  is

the number of model parameters; this assumes that the root decomposition of the covariance matrix has been computed and cached, which incurs a cost of  $O(n^3)$ . A common strategy used by some variational inference methods is to use a mean-field approach where parameters are assumed to be independent [11]. This assumption can decrease sampling costs from  $O(n^2)$  to  $O(n)$ , where  $n$  is the number of parameters. However, Riquelme et al. [45] found that TS using such approaches often leads to poorer empirical performance.

When exact inference is not possible, approximate inference methods can be used for posterior sampling. See Russo et al. [48] for a discussion of elementary approximation methods in relation to TS (e.g. Laplace approximation for unimodal distributions). Bootstrapping [20, 41, 35] is a simple procedure that fits multiple models to resampled versions of the dataset to approximate samples from the posterior distribution; the computational burden of maintaining multiple models can be prohibitive. MCMC-based methods for approximate inference, and Hamilton Monte Carlo (HMC) [40] in particular, are largely regarded as the “gold standard” for approximate Bayesian inference. HMC, and other MCMC-like approaches (e.g., Chen et al. [16], Welling and Teh [59]) generate an arbitrary number of posterior samples for all parameters. While such algorithms permit rapid evaluation of posterior samples (since the parameters are already sampled), they require substantial memory to store multiple samples of the parameters. Recent methods have also considered decomposing the covariance or precision matrix into a diagonal and low-rank component [66, 36]. While this reduces computational complexity and memory costs relative to using the full covariance, sampling still incurs a time complexity of  $O((n + 1)\rho)$  where  $\rho$  is the rank of the covariance (or precision matrix) and  $\rho$  copies of the weights must be stored.

By pre-computing and distilling TS, our imitation learning framework allows the use of the most appropriate inferential procedure for the task at hand, rather than what is feasible to run in an online setting. In particular, the separation of online decision-making and offline computation allows the use of state-of-the-art methods from the active research on deep Bayesian methods.

In contrast with previous works that consider one-shot, offline policy optimization [53, 29], our method explicitly balances exploration and exploitation across multiple horizons, while using large batch updates.

## 2 Thompson Sampling (TS)

We consider a (Bayesian) contextual bandit problem where the agent (decision-maker) observes a context, takes an action, and receives a reward. Let  $P$  be a prior distribution over a parameter space  $\Theta$ . At each time  $t$ , we denote the context  $S_t \stackrel{\text{iid}}{\sim} \mathbb{P}_S$ , action  $A_t \in \mathcal{A}$ , and reward  $R_t \in \mathbb{R}$ . We consider a well-specified reward model class  $\{f_\theta : \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R} \mid \theta \in \Theta\}$  such that  $f_\theta(a, s) = \mathbb{E}[R_t \mid \theta, A_t = a, S_t = s]$  for all  $a \in \mathcal{A}, s \in \mathcal{S}$ . Let  $H_t = (S_1, A_1, R_1, \dots, S_{t-1}, A_{t-1}, R_{t-1})$  be the history of previous observations at time  $t$ , and assume regardless of  $H_t$ , the mean reward at time  $t$  is determined only by the context-action pair  $\mathbb{E}[R_t \mid \theta, H_t, S_t = s, A_t = a] = f_\theta(a, s)$ , or equivalently,  $R_t = f_\theta(A_t, S_t) + \epsilon_t$  where  $\epsilon_t$  is a mean zero i.i.d. noise.

We use  $\pi_t$  to denote the policy at time  $t$  that generates the actions  $A_t$ , based on the history  $H_t$ : conditional on the history  $H_t$ , we have  $A_t \mid S_t \sim \pi_t(\cdot \mid S_t)$ , where we abuse notation to suppress the dependence of  $\pi_t$  on  $H_t$ . The agent’s objective is to maximize the cumulative sum of rewards by sequentially updating the policy  $\pi_t$ ’s based on observed context-action-reward tuples. The *regret* of the agent compares the agent’s cumulative rewards to the rewards under the optimal actions; for any fixed parameter value  $\theta \in \Theta$ , the (frequentist) regret for the set of policies  $\{\pi_t\}_{t \in \mathbb{N}}$  is

$$\text{Regret}(T, \{\pi_t\}_{t \in \mathbb{N}}, \theta) := \sum_{t=1}^T \mathbb{E} \left[ \max_{a \in \mathcal{A}} f_\theta(a, S_t) - f_\theta(A_t, S_t) \mid \theta \right].$$

For simplicity, we assume  $\text{argmax}_{a \in \mathcal{A}} f_\theta(a, s)$  is nonempty almost surely. We assume the agent’s prior,  $P$ , is *well-specified*, a key (standard) assumption that drives our subsequent analysis. Under the prior  $P$  over  $\theta \in \Theta$ , the Bayes regret is simply the frequentist regret averaged over  $\theta \sim P$ :

$$\text{BayesRegret}(T, \{\pi_t\}_{t \in \mathbb{N}}) := \mathbb{E}_{\theta \sim P} [\text{Regret}(T, \{\pi_t\}_{t \in \mathbb{N}}, \theta)] = \sum_{t=1}^T \mathbb{E}_{\theta \sim P} \left[ \max_{a \in \mathcal{A}} f_\theta(a, S_t) - f_\theta(A_t, S_t) \right].$$

Based on the history so far, a TS algorithm plays an action according to the posterior probability of the action being optimal. The posterior probabilities are computed based on the prior  $P$  and previously observed context-action-reward tuples. At time  $t$ , this is often implemented by sampling

---

**Algorithm 1** Imitating Thompson Sampling

---

- 1: Input: prior  $P$  on parameter space  $\Theta$ , reward model class  $\{f_\theta(\cdot, \cdot)\}$ , imitation policy model class  $\{\pi^m : m \in \mathcal{M}\}$ , notion of distance  $D$  for probabilities
  - 2: **for**  $t = 1$  **to**  $T$  **do**
  - 3:   Update  $\pi_t^{TS} | H_t$  and  $\pi_t^m$ , where  $m \leftarrow \operatorname{argmin}_{m \in \mathcal{M}} \mathbb{E}_{S \sim \mathbb{P}_S} [D(\pi_t^{TS}, \pi^m | S)]$
  - 4:   Observe  $S_t$ , sample  $A_t \sim \pi_t^m(\cdot | S_t)$ , receive  $R_t$
  - 5: **end for**
- 

from the posterior  $\theta_t \sim P(\theta \in \cdot | H_t, S_t)$ , and setting  $A_t^{TS} \in \operatorname{argmax}_{a \in \mathcal{A}} f_{\theta_t}(a, S_t)$ . By definition, TS enjoys the optimality property  $A_t^{TS} | H_t, S_t \stackrel{d}{=} A_t^* | H_t, S_t$  where  $A_t^* \in \operatorname{argmax}_{a \in \mathcal{A}} f_\theta(a, S_t)$  and  $\theta$  is the true parameter drawn from the prior  $P$ .

### 3 Imitation Learning

Motivated by challenges in implementing Thompson sampling real-time in production systems, we develop an imitation learning algorithm that separates action generation from computationally intensive steps like posterior sampling and optimization. Our algorithm maintains an explicit policy representation that emulates the TS policy by simulating its actions *offline*. At decision time, the algorithm generates an action simply by sampling from the current policy representation, which is straightforward to implement and computationally efficient to run real-time.

At each time  $t$ , our algorithm observes a context  $S_t$ , and plays an action drawn from its explicit policy representation. Formally, we parameterize our policy  $\pi^m(a | s)$  with a model class  $\mathcal{M}$  (e.g. a neural network that takes as input a context and outputs a distribution over actions) and generate actions by sampling from the current policy  $A_t \sim \pi_t^m(\cdot | S_t)$ . Upon receiving a corresponding reward  $R_t$ , the agent uses the context-action-reward tuple to update its posterior on the parameter  $\theta \in \Theta$  *offline*. Although this step requires posterior inference that may be too burdensome to run real-time, our method allows running it offline on a different computing node, so that it does not affect latency. Using the updated posterior  $\theta_t \sim \mathbb{P}(\cdot | H_t)$ , the agent then simulates actions drawn by the TS policy by computing the maximizer  $A_t^{TS}(s) \in \operatorname{argmax}_{a \in \mathcal{A}} f_{\theta_t}(a, s)$ , for a range of values  $s \in \mathcal{S}$ . Using these simulated context-action pairs, we learn an explicit policy representation that *imitates* the observed actions of the TS policy.

We summarize an idealized form of our method in Algorithm 1, where conditional on the history  $H_t$  generated by the imitation policy, we denote the off-policy TS policy at time  $t$  as  $\pi_t^{TS}(a | s)$ . This policy is different from the true, on-policy TS since the imitation policy generates actions. Dropping the subscript  $t$  to simplify notation, the imitation learning problem that updates the agent's policy is

$$\operatorname{minimize}_{m \in \mathcal{M}} \mathbb{E}_{S \sim \mathbb{P}_S} [D(\pi^{TS}, \pi^m | S)]. \quad (1)$$

Imitation problem (1) finds a model  $m \in \mathcal{M}$  that minimizes a measure of discrepancy  $D(\cdot, \cdot | S)$  between the two distributions on  $\mathcal{A}$ , conditional on the context  $S$ . This imitation objective (1) cannot be computed analytically, and we provide efficient approximation algorithms below.

To instantiate Algorithm 1, we fix Kullback-Leibler (KL) divergence as a notion of discrepancy between probabilities and present finite-sample approximations based on observed contexts and simulated actions from the TS policy  $\pi_t^{TS}(a | s)$ . For probabilities  $q^1$  and  $q^2$  on  $\mathcal{A}$  such that  $q^1, q^2 \ll \nu$  for some  $\sigma$ -finite measure  $\nu$  on  $\mathcal{A}$ , the KL divergence between  $q^1$  and  $q^2$  is  $D_{\text{kl}}(q^1 \| q^2) := \int_{\mathcal{A}} \log \frac{dq^1/d\nu}{dq^2/d\nu}(a) dP(a)$ , where we use  $dq^1/d\nu$  and  $dq^2/d\nu$  to denote Radon-Nikodym derivatives of  $q^1$  and  $q^2$  with respect to  $\nu$ . For policies  $\pi^1$  and  $\pi^2$ , let  $D_{\text{kl}}(\pi^1, \pi^2 | S) := D_{\text{kl}}(\pi^1(\cdot | S) \| \pi^2(\cdot | S))$ , where we use  $\pi^1, \pi^2$  to also denote their densities over  $\mathcal{A}$ . From the preceding display, the imitation problem (1) with  $D(\cdot, \cdot | S) = D_{\text{kl}}(\cdot, \cdot | S)$  is equivalent to a maximum log likelihood problem

$$\operatorname{maximize}_{m \in \mathcal{M}} \mathbb{E}_{S \sim \mathbb{P}_S, A^{TS} \sim \pi^{TS}(\cdot | S)} [\log \pi^m(A^{TS} | S)]. \quad (2)$$

We write  $\mathbb{E}[\cdot] = \mathbb{E}_{S \sim \mathbb{P}_S, A^{TS} \sim \pi^{TS}(\cdot | S)}[\cdot]$  for simplicity. The data comprises of context-action pairs; contexts are generated under the marginal distribution  $S \sim \mathbb{P}_S$  independent of everything else, and actions are generated/simulated from the TS policy  $A^{TS} \sim \pi^{TS}(\cdot | S)$ . The MLE problem (2) finds the model  $m \in \mathcal{M}$  that maximizes the log likelihood of observing these actions.

The imitation objective  $m \mapsto \mathbb{E}[\log \pi^m(A^{\text{TS}} \mid S)]$  involves an expectation over the unknown marginal distribution of contexts  $\mathbb{P}_S$  and actions generated by the TS policy  $\pi^{\text{TS}}(\cdot \mid S)$ . Although the expectation over  $S \sim \mathbb{P}_S$  involves a potentially high-dimensional integral over an unknown distribution, sampling from this distribution is usually very cheap since the observations  $S \sim \mathbb{P}_S$  can be “unlabeled” in the sense that no corresponding action/reward are necessary. For example, it is common for internet services to maintain a database of features  $S$  for all of its customers. Using these contexts, we can solve the MLE problem (2) efficiently via stochastic gradient descent [33, 19].

Accommodating typical application scenarios, Algorithm 1 and its empirical approximation applies to settings where the agent has the ability to interact with the system concurrently. This is a practically important feature of the algorithm, since most applications require the agent to concurrently generate actions, often with large batch sizes. Using the imitation policy  $\pi^m(a \mid s)$ , it is trivial to parallelize action generation over multiple computing nodes, even on mobile devices. Although our theoretical developments focus on the non-concurrent case for ease of exposition, our regret bounds can be extended to the batch setting; our experiments in Section 5 present large batch scenarios to illustrate typical application scenarios. While we restrict discussion to TS in this work, offline imitation can learn a explicit policy representation of any complicated policy and allow operationalization at scale.

For continuous actions with geometry, it may be natural to allow imitation policies to have slightly different support than the Thompson sampling policy. We can instantiate Algorithm 1 with Wasserstein distances as our notion of discrepancy  $D(\cdot, \cdot \mid s)$ . The subsequent theoretical development for KL divergences has its analogue for Wasserstein distances, which we outline in Section A.

## 4 Imitation Controls Regret

In this section, we show that minimizing the KL divergence (1) between TS and the imitation policy allows control on the Bayes regret of the imitation algorithm. In this sense, the imitation learning loss (1) is a valid objective where better imitation translates to gains in decision-making performance. By controlling the imitation objective (1), our results show that Algorithm 1 and its empirical approximation enjoys the similar optimality guarantees as TS. Since the imitation policy is responsible for generating actions in Algorithm 1, the observations used to update the Thompson sampling policy are different from what the TS policy would have generated. In this sense, our imitation algorithm does not emulate the *true* TS policy, but rather simply mimics its *off-policy* variant, where the posterior updates are performed based on the history generated by the imitation policy. Our analysis shows off-policy imitation is sufficient to achieve the same optimal regret bounds available for *on-policy* TS [47]. In particular, our results guard against potential compounding of errors resulting from imitating the off-policy variant of TS.

We relate the performance of our imitation policy with that of the *off-policy* TS,  $\pi^{\text{TS}}$ , and show that  $\pi^{\text{TS}}$  admits a Bayes regret decomposition that allows us to show that it still achieves same optimal Bayes regret as the (on-policy) TS policy. This allows proving Bayes regret bounds for the imitation policy by utilizing existing proofs for bounding regret of UCB algorithms. We build on the flexible approach of Russo and Van Roy [47], and extend it to imitation policies that emulate off-policy TS.

**Regret decomposition** Since our imitation learning problem (1) approximates *off-policy* Thompson sampling, a pessimistic view is that any small deviation between the imitation and TS policy can exacerbate over time. A suboptimal sequence of actions taken by the imitation policy may deteriorate the performance of the off-policy TS policy  $\pi^{\text{TS}}$  updated based on this data. Since the imitation policy again mimics the off-policy Thompson sampler, this may lead to a negative feedback loop in the worst-case. Our analysis precludes such negative cascades when outcomes are averaged over the prior  $P$ : the Bayes regret of the imitation policy is comparable to that of the best UCB algorithm, up to the sum of expected discrepancy between the off-policy TS policy and the imitation learner at each period. Our bounds imply that each single-period approximation error does not affect the Bayes regret linearly in  $T$  as our worst-case intuition suggests, but only as a one-time approximation cost. To achieve near-optimal regret, it thus suffices at each period to control the imitation objective (1) to the off-policy TS policy  $\pi_t^{\text{TS}}$ ; in Section C, we show that the imitation objective can be efficiently optimized using approximations based on unlabeled contexts.

We connect the performance of our imitation policy to that of the off-policy Thompson sampler and in turn relate the latter method’s Bayes regret to that of the *best* UCB algorithm. Let  $U_t(\cdot; H_t, S_t) : \mathcal{A} \rightarrow \mathbb{R}$  be a sequence of upper confidence bounds (UCB). Russo and Van Roy [47] showed that a

TS algorithm admits a Bayes regret decomposition with respect to *any* UCB sequence. This allows leveraging arguments that bound the regret of a UCB algorithm to bound the Bayes regret of TS. Since the Bayes regret decomposition for TS holds for *any* UCB sequence, the performance of TS enjoys Bayes regret guarantees of the best UCB algorithm. We can show a similar Bayes regret decomposition for our imitation policy for any UCB sequence; the Bayes regret of the imitation policy enjoys a UCB regret decomposition similar to TS, up to the cumulative sum of single-period approximation errors. Recall that we denote  $A_t^{\text{TS}} \sim \pi_t^{\text{TS}}(\cdot | S_t)$ , the action generated by the off-policy Thompson sampler. See Section D.1 for the proof of the following result.

**Lemma 1.** *Let  $\{\pi_t\}_{t \in \mathbb{N}}$  be any sequence of policies (adapted to history), and let  $U_t(\cdot; H_t, S_t) : \mathcal{A} \rightarrow \mathbb{R}$  be any upper confidence bound sequence that is measurable with respect to  $(H_t, S_t)$ . Let there be a sequence  $M_t(H_t, S_t)$  and  $L > 0$  such that  $\sqrt{\mathbb{E}[M_t(H_t, S_t)^2]} \leq L$ , and  $\sup_{a \in \mathcal{A}} |U_t(a; H_t, S_t)| \leq M_t(H_t, S_t)$ . Then, for all  $T \in \mathbb{N}$ ,  $\text{BayesRegret}(T, \{\pi_t\}_{t \in \mathbb{N}})$  is bounded by*

$$\sum_{t=1}^T \mathbb{E}[f_\theta(A_t^*, S_t) - U_t(A_t^*; H_t, S_t)] + \sum_{t=1}^T \mathbb{E}[U_t(A_t; H_t, S_t) - f_\theta(A_t, S_t)] + L \sum_{t=1}^T \sqrt{2\mathbb{E}[D_{\text{kl}}(\pi_t^{\text{TS}}, \pi_t | S_t)]}.$$

The above Bayes regret decomposition shows that performance analysis of any UCB algorithm can characterize the regret of our imitation policy. In this sense, the imitation policy achieves regret comparable to the *optimal* UCB algorithm, up to the sum of single-period imitation errors. The first two terms in the regret decompositions can be bounded using canonical UCB proofs following the approach of Russo and Van Roy [47]. We detail such arguments in concrete modeling scenarios in the next subsection, where our results draw from regret guarantees for UCB algorithms [1, 2, 51]. The last term in the decomposition is controlled by our imitation learning algorithm (Algorithm 1) and its empirical approximation; we show in Section C that the cheap availability of unlabeled contexts allows tight control of imitation errors, with a near-optimal dimension dependence.

The fact that we are studying Bayes regret, as opposed to the frequentist regret, plays an important role in the above decomposition. We conjecture that in the worst-case, even initially small imitation error (and consequently suboptimal exploration) can compound over time in a linear manner. It remains open whether certain problem structures can provably preclude these negative feedback cycles uniformly over  $\theta$ , which is necessary for obtaining frequentist regret bounds under imitation.

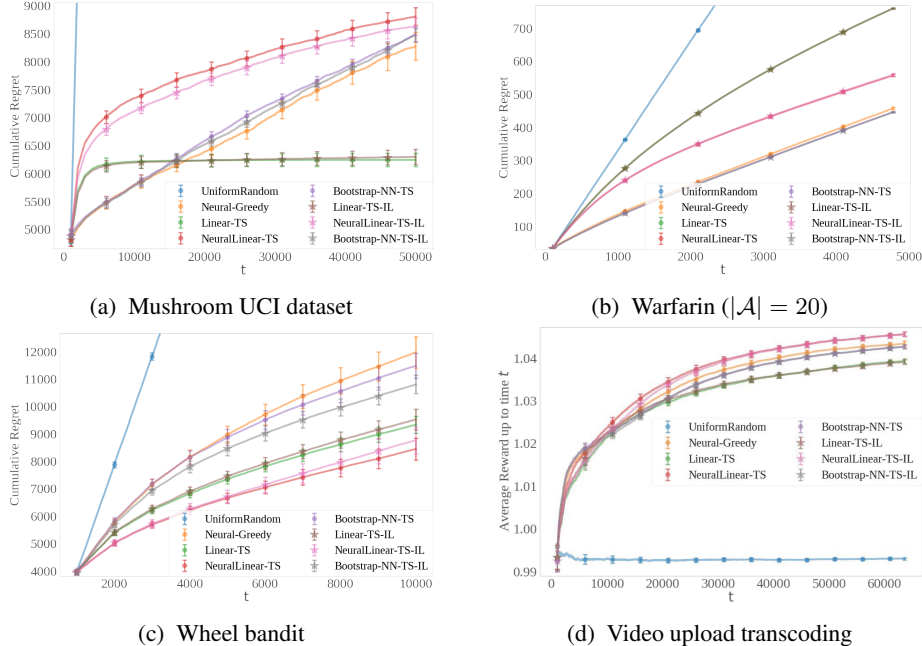
**Regret bounds** Building on the Bayes regret decomposition, we now show concrete regret guarantees for our imitation algorithm. By leveraging analysis of UCB algorithms, we proceed by bounding the first two sums in the decomposition. Our bounds on the Bayes regret are instance-independent (gap-independent), and shows that the imitation policy achieves optimal regret (or best known bounds thereof) up to the sum of imitation error terms. We present two modeling scenarios, focusing first on the setting where the mean reward function  $(a, s) \mapsto f_\theta(a, s)$  can be modeled as a generalized linear model. Secondly, we consider the modeling the mean reward function  $(a, s) \mapsto f_\theta(a, s)$  as a Gaussian process in Section B; in this nonparametric setting, our regret bounds scale with the maximal information gain possible over  $T$  rounds.

Let the mean reward function  $f_\theta(a, s)$  be modeled as a generalized linear model. We assume  $\Theta \subseteq \mathbb{R}^d$ , and that there exists a feature vector  $\phi : \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$  and a link function  $g : \mathbb{R} \rightarrow \mathbb{R}$  satisfying  $f_\theta(a, s) = g(\theta^\top \phi(a, s))$  for all  $\theta \in \Theta, a \in \mathcal{A}, s \in \mathcal{S}$ . Armed with the regret decomposition in Lemma 1, we obtain regret bounds by following an eluder dimension argument pioneered by Russo and Van Roy [47]. We give its proof in Section D.2.

**Theorem 1.** *Let  $g : \mathbb{R} \rightarrow \mathbb{R}$  and  $\phi : \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^d$  such that  $f_\theta(a, s) = g(\phi(a, s)^\top \theta)$  for all  $\theta \in \Theta$ , where  $g$  is an increasing, differentiable, 1-Lipschitz function. Let  $c_1, c_2, \sigma > 0$  be such that  $\sup_{\theta \in \Theta} \|\theta\|_2 \leq c_1$ , and  $\sup_{a \in \mathcal{A}, s \in \mathcal{S}} \|\phi(a, s)\|_2 \leq c_2$ , and assume that  $R_t - f_\theta(A_t, S_t)$  is  $\sigma$ -sub-Gaussian conditional on  $(\theta, H_t, S_t, A_t)$ . If  $r$  is the maximal ratio of the slope of  $g$   $r := \frac{\sup_{\theta \in \Theta, a \in \mathcal{A}, s \in \mathcal{S}} g'(\phi(a, s)^\top \theta)}{\inf_{\theta \in \Theta, a \in \mathcal{A}, s \in \mathcal{S}} g'(\phi(a, s)^\top \theta)}$ , then there is a constant  $C$  that depends on  $c_1, c_2$  such that*

$$\text{BayesRegret}(T, \{\pi_t\}_{t \in \mathbb{N}}) \leq C(\sigma + 1)rd\sqrt{T} \log rT + c_1 c_2 \sum_{t=1}^T \sqrt{2\mathbb{E}[D_{\text{kl}}(\pi_t^{\text{TS}}, \pi_t | S_t)]}. \quad (3)$$

For linear contextual bandits  $g(x) = x$ , our upper bound on the Bayes regret is tight up to a factor of  $\log T$  and the cumulative sum of the imitation errors [46]; for generalized linear models, the first two terms are the tightest bounds on the regret available. We conclude that controlling the imitation error directly controls Bayes regret. As we present in Section C, solving the empirical approximation to the imitation problem (1) with a sufficiently rich imitation model class  $m \in \mathcal{M}$  guarantees low imitation errors in typical application scenarios.



**Figure 1.** For all benchmark problems, we report the mean cumulative regret (or running average regret up to time  $t$  for the video upload transcoding application) and two standard errors of the mean over 50 trials (except for the Wheel which we run for 100 trials due the rare large rewards).

	MUSHROOM	WHEEL	VIDEO TRANSCODE	WARFARIN
UNIFORMRANDOM	0.040 ( $\pm 0.000$ )	0.039 ( $\pm 0.000$ )	0.040 ( $\pm 0.000$ )	0.040 ( $\pm 0.000$ )
NEURAL-GREEDY	0.242 ( $\pm 0.001$ )	0.228 ( $\pm 0.001$ )	0.231 ( $\pm 0.001$ )	0.232 ( $\pm 0.000$ )
LINEAR-TS	0.715 ( $\pm 0.001$ )	1.142 ( $\pm 0.001$ )	1.575 ( $\pm 0.002$ )	3.963 ( $\pm 0.002$ )
NEURALLINEAR-TS	0.826 ( $\pm 0.001$ )	1.492 ( $\pm 0.001$ )	1.931 ( $\pm 0.002$ )	4.814 ( $\pm 0.004$ )
BOOTSTRAP-NN-TS	0.235 ( $\pm 0.001$ )	0.235 ( $\pm 0.001$ )	0.236 ( $\pm 0.001$ )	0.226 ( $\pm 0.001$ )
LINEAR-TS-IL	0.184 ( $\pm 0.001$ )	0.178 ( $\pm 0.000$ )	0.169 ( $\pm 0.000$ )	0.175 ( $\pm 0.000$ )
NEURALLINEAR-TS-IL	0.186 ( $\pm 0.000$ )	0.179 ( $\pm 0.001$ )	0.169 ( $\pm 0.000$ )	0.175 ( $\pm 0.000$ )
BOOTSTRAP-NN-TS-IL	0.190 ( $\pm 0.001$ )	0.178 ( $\pm 0.000$ )	0.175 ( $\pm 0.000$ )	0.179 ( $\pm 0.001$ )

**Table 1.** Decision-making latency in milliseconds. All latency measurements were made on a Intel Xeon E5-2680 v4 @ 2.40GHz CPU with 32-bit floating point precision. For each latency measurement, action generation is repeated 100K times and the mean latency and its 2-standard errors are reported.

## 5 Empirical Evaluation and Discussion

We now empirically demonstrate our imitation algorithm on real and simulated problems. As our main real-world application, we consider a internet service receiving millions of video upload requests per day. We learn policies that decide how to transcode a video at upload time, where latency of decisions is important to the quality of service and keeping the user engaged. We observe that our IL algorithm achieves regret comparable to that of TS, while significantly reducing latency on all problems. We begin by providing a brief description of benchmark problems, deferring their details to Section F.3.

**MUSHROOM** is a UCI dataset [18] containing 8,124 examples with 22 categorical features about the mushroom and binary labels (poisonous or not). The agent decides whether to eat the mushroom or not and receives a stochastic reward: with equal probability, eating a poisonous mushroom lead to illness ( $R_t = -35$ ) or no harm ( $R_t = 5$ ), while a nonpoisonous mushroom is always harmless ( $R_t = 5$ ). The reward for abstaining is always  $R_t = 0$ . We use 50,000 contexts for each trial. **WARFARIN** is a pharmacological dosage optimization example, where we learn optimal personalized dosages for Warfarin, the most widely used anticoagulant in the US. The dataset [60, 65] contains optimal dosages for 4,788 patients, alongside 17-dimensional genetic and demographic features. We construct a CB problem with 20 discrete dosage levels as actions, where the reward is the difference between the patient’s prescribed and optimal dosage. We reshuffle contexts for each trial; similar results hold when 50,000 contexts are re-sampled each trial and 50 actions are used (Appendix

F). As our third benchmark, to evaluate our methods on examples where exploration matters, we consider the **Wheel** problem, a 2-dimensional synthetic problem specifically designed to emphasize exploration [45]. There are 5 actions and rarely seen contexts yield high rewards under one context-dependent action. We sample 10,000 contexts for each trial. As our last experiment, we consider a real-world video upload transcoding optimization problem (**VIDEO TRANSCODING**). This logged dataset contains 8M observations of video upload transcoding decisions collected by Facebook under a uniform random policy. The 38-dimensional contexts contain information about the video and network connection, and 7 actions corresponding to different video codec qualities. The rewards for successful uploads are positive and are monotonically increasing with quality. The reward is zero if the upload fails. We evaluate the performance of different contextual bandit algorithms using the rejection sampling technique proposed by Li et al. [34]. We sample 50,000 contexts for each trial.

We consider a variety of models among those reported by Riquelme et al. [45] to perform the best in a broad range of benchmark problems. The hyperparameters for the Thompson sampling algorithms are from Riquelme et al. [45], which we detail in Appendix F.2 along with those for the imitation algorithm. Policies are updated every 1000 contexts (except for on the Warfarin problem, where we update policies every 100 contexts due to the small size of the dataset) and are initialized using a uniform random policy before the first batch update. **LINEAR-TS** uses an exact Bayesian linear regression to model the reward distribution for each action  $a$ . Exact posterior samples are used under the assumption that the data for action  $a$  were generated from the linear function:  $r_a = \mathbf{s}^T \boldsymbol{\theta}_a + \varepsilon$  where  $\varepsilon \sim \mathcal{N}(0, \sigma_a^2)$ . For each action, we independently model the joint distribution,  $P(\boldsymbol{\theta}, \sigma^2) = P(\boldsymbol{\theta}|\sigma^2)P(\sigma^2)$  as a normal-inverse-gamma distribution which allows for tractable posterior inference (see Appendix F.2 for details). **NEURALLINEAR-TS** models rewards using a neural network with two 100-unit hidden layers and ReLU activations, and uses the last hidden layer  $\phi(\mathbf{s})$  as the feature representation for a Linear-TS policy. The neural network takes the context as input and has an output for each action corresponding to the reward under each action. The parameters of the neural network are shared for all actions and are learned independently of the Bayesian linear models. **BOOTSTRAP-NN-TS** trains multiple neural networks on bootstrapped observations and randomly samples a single network to use for each decision. **TS-IL** imitates a TS policy using a fully-connected neural network as its policy representation. The policy network has two 100-unit hidden layers and tanh activations, and a final soft-max layer. We compare TS and its imitation counterparts against two additional methods: a random policy (**UNIFORMRANDOM**) and a greedy policy that uses a feed-forward neural network (**NEURAL-GREEDY**).

Figure 1 shows that each TS-IL method achieves comparable performance to its corresponding vanilla TS algorithm on each respective benchmark problem. In terms of cumulative performance, each TS-IL policy consistently matches its corresponding TS policy over time. For all methods, we also evaluate decision-time latency—the time required for a policy to select an action when it is queried—and time complexity for online action-generation. While **BOOTSTRAP-NN-TS** achieves low prediction latency, it requires storing many replicates of the neural network and can significantly increase the memory footprint. Table 1 shows that the imitation policies (TS-IL) have significantly lower decision time latency compared to TS algorithms, often by over an order of magnitude on problems with larger action spaces (Warfarin and video upload transcoding). This is because generating an action under the vanilla TS policies requires drawing a sample from the joint posterior  $P(\boldsymbol{\theta}_a, \sigma_a^2)$  for each of the actions  $a$ , which is quadratic with respect to the context dimension for **LINEAR-TS** or the size of the last hidden layer for **NEURALLINEAR-TS**. TS-IL simply requires a forward pass through policy network and a cheap multinomial sample. Latency and complexity may be even greater under inference schemes not considered here (see Appendix G for a discussion).

**Discussion** In this paper, we demonstrated that Thompson sampling via imitation learning is a simple, practical, and efficient method for batch Thompson sampling with desirable regret properties. By distilling the Thompson sampling policy into easy-to-deploy explicit policy representations (e.g. small neural networks), we allow state-of-the-art Bayesian approaches to be used in contextual bandit problems. We hope that this work facilitates applications of modern Bayesian approaches in large-scale contextual bandit problems.



## References

- [1] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- [2] Y. Abbasi-Yadkori, D. Pal, and C. Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9, 2012.
- [3] M. Abeille, A. Lazaric, et al. Linear thompson sampling revisited. *Electronic Journal of Statistics*, 11(2):5165–5197, 2017.
- [4] R. J. Adler and J. E. Taylor. *Random fields and geometry*, volume 115. Springer, 2009.
- [5] D. Agarwal, B. Long, J. Traupman, D. Xin, and L. Zhang. Laser: A scalable response prediction platform for online advertising. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 173–182. ACM, 2014.
- [6] S. Agrawal and N. Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proceedings of the Twenty Fifth Annual Conference on Computational Learning Theory*, 2012.
- [7] S. Agrawal and N. Goyal. Further optimal regret bounds for thompson sampling. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, pages 99–107, 2013.
- [8] S. Agrawal and N. Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 127–135, 2013.
- [9] P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [10] K. Bhardwaj, C.-Y. Lin, A. Sartor, and R. Marculescu. Memory- and communication-aware model compression for distributed deep learning inference on iot. *ACM Trans. Embed. Comput. Syst.*, 18(5s), Oct. 2019. doi: 10.1145/3358205.
- [11] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622, 2015.
- [12] J. F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems*. Springer, 2000.
- [13] S. Boucheron, O. Bousquet, and G. Lugosi. Theory of classification: a survey of some recent advances. *ESAIM: Probability and Statistics*, 9:323–375, 2005.
- [14] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: a Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- [15] O. Chapelle and L. Li. An empirical evaluation of thompson sampling. In *Advances in Neural Information Processing Systems 24*, pages 2249–2257, 2011.
- [16] T. Chen, E. Fox, and C. Guestrin. Stochastic gradient hamiltonian monte carlo. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, 2014.
- [17] V. Dani, T. Hayes, and S. Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the Twenty First Annual Conference on Computational Learning Theory*, 2008.
- [18] D. Dua and C. Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- [19] J. C. Duchi. Introductory lectures on stochastic convex optimization. In *The Mathematics of Data*, IAS/Park City Mathematics Series. American Mathematical Society, 2018.
- [20] D. Eckles and M. Kaptein. Thompson sampling with the online bootstrap. *arXiv preprint arXiv:1410.4009*, 2014.
- [21] K. J. Ferreira, D. Simchi-Levi, and H. Wang. Online network revenue management using thompson sampling. *Operations Research*, 66(6):1586–1602, 2018.

- [22] S. Fujimoto, D. Meger, and D. Precup. Off-policy deep reinforcement learning without exploration. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2052–2062, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- [23] J. Gauci, E. Conti, Y. Liang, K. Virochsiri, Y. He, Z. Kaden, V. Narayanan, X. Ye, Z. Chen, and S. Fujimoto. Horizon: Facebook’s open source applied reinforcement learning platform. *arXiv:1811.00260 [cs.LG]*, 2018.
- [24] S. Ghosal, A. Roy, et al. Posterior consistency of gaussian process prior for nonparametric binary regression. *Annals of Statistics*, 34(5):2413–2429, 2006.
- [25] A. Gopalan, S. Mannor, and Y. Mansour. Thompson sampling for complex online problems. In *Proceedings of the 31st International Conference on Machine Learning*, pages 100–108, 2014.
- [26] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.
- [27] D. N. Hill, H. Nassif, Y. Liu, A. Iyer, and S. Vishwanathan. An efficient bandit algorithm for realtime multivariate optimization. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Aug 2017. doi: 10.1145/3097983.3098184. URL <http://dx.doi.org/10.1145/3097983.3098184>.
- [28] J. Honda and A. Takemura. Optimality of thompson sampling for gaussian bandits depends on priors. In *Artificial Intelligence and Statistics*, pages 375–383, 2014.
- [29] T. Joachims, A. Swaminathan, and M. de Rijke. Deep learning with logged bandit feedback. In *International Conference on Learning Representations*, 2018. URL [https://openreview.net/forum?id=SJaP\\_-xAb](https://openreview.net/forum?id=SJaP_-xAb).
- [30] E. Kaufmann, N. Korda, and R. Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *International Conference on Algorithmic Learning Theory*, pages 199–213. Springer, 2012.
- [31] J. Kawale, H. H. Bui, B. Kveton, L. Tran-Thanh, and S. Chawla. Efficient thompson sampling for online matrix factorization recommendation. In *Advances in Neural Information Processing Systems 15*, pages 1297–1305, 2015.
- [32] A. Krause and C. S. Ong. Contextual gaussian process bandit optimization. In *Advances in Neural Information Processing Systems 24*, pages 2447–2455, 2011.
- [33] H. J. Kushner and G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, second edition, 2003.
- [34] L. Li, W. Chu, J. Langford, and X. Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM ’11*, page 297–306. Association for Computing Machinery, 2011. doi: 10.1145/1935826.1935878.
- [35] X. Lu and B. Van Roy. Ensemble sampling. In *Advances in neural information processing systems*, pages 3258–3266, 2017.
- [36] W. J. Maddox, P. Izmailov, T. Garipov, D. P. Vetrov, and A. G. Wilson. A simple baseline for bayesian uncertainty in deep learning. In *Advances in Neural Information Processing Systems 32*, 2019.
- [37] H. Mao, S. Chen, D. Dimmery, S. Singh, D. Blaisdell, Y. Tian, M. Alizadeh, and E. Bakshy. Real-world video adaptation with reinforcement learning. In *Neural Information Processing Systems Workshop on RL for Real Life*, 2019.
- [38] A. Mas-Colell, M. D. Whinston, J. R. Green, et al. *Microeconomic theory*, volume 1. Oxford university press New York, 1995.

- [39] B. C. May and D. S. Leslie. Simulation studies in optimistic bayesian sampling in contextual-bandit problems. *Technical Report, Statistics Group, Department of Mathematics, University of Bristol*, 11:02, 2011.
- [40] R. M. Neal. Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2, 2011.
- [41] I. Osband, C. Blundell, A. Pritzel, and B. Van Roy. Deep exploration via bootstrapped dqn. In *Advances in neural information processing systems*, pages 4026–4034, 2016.
- [42] A. Petrescu and S. Tas. Client side ranking to more efficiently show people stories in feed, Oct 2016. URL <https://engineering.fb.com/networking-traffic/client-side-ranking-to-more-efficiently-show-people-stories-in-feed/>.
- [43] G. Peyré, M. Cuturi, et al. Computational optimal transport. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [44] D. Pollard. Empirical processes: theory and applications. In *NSF-CBMS regional conference series in probability and statistics*, pages i–86. JSTOR, 1990.
- [45] C. Riquelme, G. Tucker, and J. Snoek. Deep bayesian bandits showdown. In *International Conference on Learning Representations*, 2018.
- [46] P. Rusmevichientong and J. N. Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.
- [47] D. Russo and B. Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- [48] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen. A tutorial on thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018.
- [49] E. M. Schwartz, E. T. Bradlow, and P. S. Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017.
- [50] S. L. Scott. A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26(6):639–658, 2010.
- [51] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.
- [52] Statista. Percentage of all global web pages served to mobile phones from 2009 to 2018, 2020. URL <https://www-statista-com/statistics/241462/global-mobile-phone-website-traffic-share/>.
- [53] A. Swaminathan and T. Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *Journal of Machine Learning Research*, 16(52):1731–1755, 2015. URL <http://jmlr.org/papers/v16/swaminathan15a.html>.
- [54] W. 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.
- [55] A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009.
- [56] A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer, New York, 1996.
- [57] C. Villani. *Optimal Transport: Old and New*. Springer, 2009.
- [58] M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.
- [59] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688, 2011.

- [60] M. Whirl-Carrillo, E. McDonagh, J. Hebert, I. Gong, K. Sangkuhl, C. Thorn, R. Altman, and T. Klein. Pharmacogenomics knowledge for personalized medicine. *Clinical pharmacology and therapeutics*, 92:414–7, 10 2012. doi: 10.1038/clpt.2012.96.
- [61] A. Wilson and H. Nickisch. Kernel interpolation for scalable structured gaussian processes (kiss-gp). In *International Conference on Machine Learning*, pages 1775–1784, 2015.
- [62] A. G. Wilson, C. Dann, and H. Nickisch. Thoughts on massively scalable gaussian processes. *arXiv:1511.01870 [cs.LG]*, 2015.
- [63] A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing. Deep kernel learning. In *Artificial Intelligence and Statistics*, pages 370–378, 2016.
- [64] C.-J. Wu, D. Brooks, K. Chen, D. Chen, S. Choudhury, M. Dukhan, K. Hazelwood, E. Isaac, Y. Jia, B. Jia, T. Leyvand, H. Lu, Y. Lu, L. Qiao, B. Reagen, J. Spisak, F. Sun, A. Tulloch, P. Vajda, X. Wang, Y. Wang, B. Wasti, Y. Wu, R. Xian, S. Yoo, and P. Zhang. Machine learning at facebook: Understanding inference at the edge. In *Proceedings - 25th IEEE International Symposium on High Performance Computer Architecture, HPCA 2019*, pages 331–344, 3 2019. doi: 10.1109/HPCA.2019.00048.
- [65] H. Xiao. Online learning to estimate warfarin dose with contextual linear bandits. *CoRR*, 2019.
- [66] G. Zhang, S. Sun, D. Duvenaud, and R. Grosse. Noisy natural gradient as variational inference. In *International Conference on Machine Learning*, pages 5852–5861, 2018.