

Toward Optimal Solution for the Context-Attentive Bandit Problem

Djallel Bouneffouf¹, Raphael Feraud², Sohini Upadhyay¹, Irina Rish³ and Yasaman Khazaeni¹

¹IBM Thomas J. Watson Research Center, Yorktown Heights, NY USA

²Orange Labs

³Universite de montreal

¹{Djallel.Bouneffouf@, Sohini.upadhyay@, Yasaman.khazaeni@}ibm.com,

²Raphael.Feraud@orange.com, ³irina.rish@mila.quebec

Abstract

In various recommender system applications, from medical diagnosis to dialog systems, due to observation costs only a small subset of a potentially large number of context variables can be observed at each iteration; however, the agent has a freedom to choose which variables to observe. In this paper, we analyze and extend an online learning framework known as *Context-Attentive Bandit*. We derive a novel algorithm, called *Context-Attentive Thompson Sampling (CATS)*, which builds upon the Linear Thompson Sampling approach, adapting it to *Context-Attentive Bandit* setting. We provide a theoretical regret analysis and an extensive empirical evaluation demonstrating advantages of the proposed approach over several baseline methods on a variety of real-life datasets.

1 Introduction

The contextual bandit (CB) problem has been extensively studied in the past, and a variety of solutions have been proposed. In LINUCB [Li *et al.*, 2010; Abbasi-Yadkori *et al.*, 2011; Bouneffouf *et al.*, 2019; Bouneffouf and Rish, 2019], Neural Bandit [Allesiardo *et al.*, 2014] and in linear Thompson Sampling [Agrawal and Goyal, 2013; Balakrishnan *et al.*, 2019a; Balakrishnan *et al.*, 2019b], a linear dependency is assumed between the expected reward given the context and an action taken after observing this context; the representation space is modeled using a set of linear predictors.

Recently, a promising variant of contextual bandits, called a *Context Attentive Bandit (CAB)* was proposed in [Bouneffouf *et al.*, 2017], where no context is given by default, but the agent can request to observe (to focus its "attention" on) a limited number of context variables at each iteration. We propose here an extension of this problem setting: a small subset of V context variables is revealed at each iteration (i.e. partially observable context), followed by the agent's choice of additional U features, where V , U , and the set of V immediately observed features are fixed at all iterations. The agent must learn to select both the best additional features and, subsequently, the best arm to play, given the resulting $V + U$ observed features. (The original *Context Attentive Bandit* corresponds to $V = 0$.)

The proposed setting is motivated by several real-life applications where observing the full context is impossible and the agent is allowed to observe just a small subset of information at the onset.

For instance, in a clinical setting, a doctor may first take a look at patient's medical record (partially observed context) to decide which medical test (additional context variables) to perform, before choosing a treatment plan (selecting an arm to play). It is often too costly or even impossible to conduct all possible tests (i.e., observe the full context); therefore, given the limit on the number of tests, the doctor must decide which subset of tests will result into maximally effective treatment choice (maximize the reward).

Similar problems can arise in multi-skill orchestration for AI agents. For example, in dialog orchestration, a user's query is first directed to a number of domain-specific agents, each providing a different response, and then the best response is selected. However, it might be too costly to request the answers from *all* domain-specific experts, especially in multi-purpose dialog systems with a very large number of domains experts. Given a limit on the number of experts to use for each query, the orchestration agent must choose the best subset of experts to use. In this application, the query is the immediately observed part of the overall context, while the responses of domain-specific experts are the initially unobserved features from which a limited subset must be selected and observed, before choosing an arm, i.e. deciding on the best out of the available responses. For multi-purpose dialog systems, such as, for example, personal home assistants, retrieving features or responses from every domain-specific agent is computationally expensive or intractable, with the potential to cause a poor user experience, again underscoring the need for effective feature selection.

Overall, the main contributions of this paper include:

- (1) a generalization of *Context Attentive Bandit*, and a first lower bound for this problem.
- (2) an algorithm called *Context Attentive Thompson Sampling* for stationary and non-stationary environments, and its regret bound in the case of stationary environment.
- (3) an extensive empirical evaluation demonstrating advantages of our proposed algorithm over the previous context-attentive bandit approach [Bouneffouf *et al.*, 2017], on a range of datasets in both stationary and non-stationary settings.

This paper is organized as follows. In section 2 we review related works. The section 3 introduces some background concepts. Then we introduce the *Context-Attentive Bandit*, the proposed algorithms for both stationary and non-stationary environments, and we provide a lower bound of the proposed problem and a regret bound of the proposed algorithm in the case of stationary environment. Experimental evaluation on several datasets, for varying parameter settings, is presented afterward. Finally, the last section concludes the paper and points out possible directions for future works.

2 Related Work

In the *contextual bandit problem* [Langford and Zhang, 2008; Agarwal *et al.*, 2009; Auer *et al.*, 2003], the objective of is to learn the relationship between the context and reward, in order to find the best arm-selection policy for maximizing cumulative reward over time. Motivated by dimensionality reduction tasks, [Abbasi-Yadkori *et al.*, 2012] studied a sparse variant of stochastic linear bandits, where only a relatively small and unknown subset of features is relevant to a multi-variate function optimization. It presents an application to the problem of optimizing a function that depends on many features, where only a small, initially unknown subset of features is relevant. Similarly, [Carpentier and Munos, 2012] also considered high-dimensional stochastic linear bandits with sparsity. There the authors combined ideas from compressed sensing and bandit theory to derive a novel algorithm. In [Oswal *et al.*, 2020], authors explore a new form of the linear bandit problem in which the algorithm receives the usual stochastic rewards as well as stochastic feedback about which features are relevant to the rewards and propose an algorithm that can achieve $O(\sqrt{T})$ regret, without prior knowledge of which features are relevant.

In [Bouneffouf *et al.*, 2017] the authors proposed the novel framework of *contextual bandit with restricted context*, where observing the whole feature vector at each iteration is too costly or impossible for some reasons; however, the agent can request to observe the values of an arbitrary subset of features within a given budget, i.e. the limit on the number of features observed. This paper explores a more general problem, and unlike [Bouneffouf *et al.*, 2017], we provide a theoretical analysis of the proposed problem and the proposed algorithm.

The *Context Attentive Bandit* problem is related to the *budgeted learning* problem, where a learner can access only a limited number of attributes from the training set or from the test set (see for instance [Cesa-Bianchi *et al.*, 2011]). In [Foster *et al.*, 2016], the authors studied the *online budgeted learning* problem. They showed a significant negative result: for any $\delta > 0$ no algorithm can achieve regret bounded by $O(T^{1-\delta})$ in polynomial time. For overcoming this negative result, an additional assumption is necessary. Here, following [Durand and Gagné, 2014], we assume that the expected reward of selecting a subset of features is the sum of the expected rewards of selecting individually the features. We obtain an efficient algorithm, which has a linear algorithmic complexity in terms of time horizon.

3 Context Attentive Bandit Problem

We now introduce the problem setting, outlined in Algorithm 1. Let C be a set of N features. At each time point t the environment generates a feature vector $\mathbf{c}(t) = (c_1(t), \dots, c_N(t), 1) \in \mathbb{R}^{N+1}$, which the agent cannot observe fully, but a partial observation of the context is allowed: the values of a subset of $V < N$ observed features $C^V \subset C$, are revealed: $\mathbf{c}^V = (c_1^V(t), \dots, c_N^V(t), 1) \in \mathbb{R}^{N+1}$, $\forall i \in C, c_i^V(t) = c_i(t)\mathbb{1}\{i \in C^V\}$. Based on this partially observed context, the agent is allowed to request an additional subset of U unobserved features $C^U \subset C \setminus C^V$, $V + U \leq N$. The goal of the agent is to maximize its total reward over time via (1) the optimal choice of the additional set of features C^U , given the initial observed features $\mathbf{c}^V(t)$, and (2) the optimal choice of an arm $k \in [K] = \{1, \dots, K\}$ based on $\mathbf{c}^{V+U}(t) = (c_1^{V+U}(t), \dots, c_N^{V+U}(t), 1) \in \mathbb{R}^{N+1}$, $\forall i \in C, c_i^{V+U}(t) = c_i(t)\mathbb{1}\{i \in C^V \vee i \in C^U\}$. We assume $P_r(r|\mathbf{c}, k)$, an unknown probability distribution of the reward given the context and the action taken in that context. In the following the expectations are taken over the probability distribution $P_r(r|\mathbf{c}, k)$.

Algorithm 1 Context Attentive Bandit Problem (CAB)

- 1: **for** $t := 1$ to T **do**
 - 2: Context $\mathbf{c}(t)$ is chosen by the environment
 - 3: The values $\mathbf{c}^V(t)$ of a subset $C^V \subset C$ are revealed
 - 4: The agent selects a subset $C^U \subseteq C \setminus C^V$
 - 5: The values $\mathbf{c}^U(t)$ are revealed;
 - 6: The agent chooses an arm $k(t) := \pi(\mathbf{c}^{V+U}(t))$
 - 7: The reward $r_{k(t)}$ is sampled from distribution $P_r(r|\mathbf{c}, k)$ and it is revealed
 - 8: The agent updates the policy $\pi \in \Pi_{C^{V+U}}$
 - 9:
 - 10: **end for**
-

The contextual bandit problem. Following [Langford and Zhang, 2008], this problem is defined as follows. At each time point $t \in \{1, \dots, T\}$, an agent is presented with a *context* (i.e. *feature vector*) $\mathbf{c}(t) \in \mathbb{R}^{N+1}$ before choosing an arm $k \in [K]$. Let $\mathbf{r}(t) = (r_1(t), \dots, r_K(t)) \in \mathbb{R}^K$ denote a reward vector, where $r_k(t)$ is the reward at time t associated with the arm k . Let $\pi : \mathbb{R}^{N+1} \rightarrow [K]$ denote a policy, mapping a context $\mathbf{c}(t) \in \mathbb{R}^{N+1}$ into an arm $k \in [K]$. We assume that the expected reward is a linear function of the context.

Assumption 1 (linear contextual bandit): *Whatever the subset of selected features $C^U \subset C \setminus C^V$ the expected reward is a linear function of the context: $\mathbb{E}[r_k|\mathbf{c}^{V+U}(t)] = \mathbf{c}^{V+U}(t)^\top \boldsymbol{\mu}_k$, where $\boldsymbol{\mu}_k \in \mathbb{R}^{N+1}$ is an unknown parameter, $\forall k \in [K] \|\boldsymbol{\mu}_k\| \leq 1$, and $\|\mathbf{c}^{V+U}(t)\| \leq 1$.*

Contextual Combinatorial Bandit. The contextual combinatorial bandit problem [Qin *et al.*, 2014] can be viewed as a game where the agent sequentially observes a context $\mathbf{c}^V(t)$, selects a subset $C^U(t) \subset C \setminus \{C^V\}$ and observes the reward corresponding to the selected subset. The goal is to maximize the reward over time. Let $r_U|\mathbf{c}^V(t), \pi \in \mathbb{R}$ be the reward associated with the set of selected features C^U knowing the context vector $\mathbf{c}^{V+U}(t)$ and the policy π . We have

$r_U | \mathbf{c}^V(t), \pi = r_{k(t)} | \mathbf{c}^{V+U}(t)$, where $k(t) = \pi(\mathbf{c}^{V+U}(t))$. Each feature $i \in C^U$ is associated with the corresponding random variable $r_i | \mathbf{c}^V(t), \pi \in \mathbb{R}$ which indicates the reward obtained when choosing the i -th feature at time t .

Assumption 2 (linear contextual combinatorial bandit): the mean reward of selecting the set of features $C^U \subset C \setminus \{C^V\}$ is: $\mathbb{E}[r_U | \mathbf{c}^V(t), \pi] = \sum_{i \in C^U} \mathbb{E}[r_i | \mathbf{c}^V(t), \pi]$, and the expectation of the reward of selecting the feature i is a linear function of the context vector $\mathbf{c}^V(t)$: $\mathbb{E}[r_i | \mathbf{c}^V(t), \pi] = \mathbf{c}^V(t)^\top \boldsymbol{\theta}_i$, where $\boldsymbol{\theta}_i \in \mathbb{R}^{N+1}$ is an unknown weight vector associated with the feature i , $\forall i \in C \setminus C^V \|\boldsymbol{\theta}_i\| \leq 1$, and $\|\mathbf{c}^V(t)\| \leq 1$. Let Π_{C^V+U} be the set of linear policies such that only the features coming from C^V+U are used, where C^V is a fixed subset of C , and C^U is any subset of $C \setminus C^V$. The objective of *Contextual Attentive Bandit* (Algorithm 1) is to find an optimal policy $\pi^* \in \Pi_{C^V+U}$, over T iterations or time points, so that the total reward is maximized.

Definition 1 (Optimal Policy for CAB). The optimal policy π^* for handling the CAB problem is selecting the arm at time t :

$$\begin{aligned} k^*(t) &= \arg \max_{k \in [K], C^U \in C \setminus C^V} \mathbf{c}^{U+V}(t)^\top \boldsymbol{\mu}_k \\ &= \arg \max_{k \in [K]} \mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_k, \end{aligned}$$

where

$$\begin{aligned} C^{U^*} &= \arg \max_{C^U \subset C \setminus C^V} r_U | \mathbf{c}^V(t), \boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K = \\ &= \arg \max_{C^U \subset C \setminus C^V} \sum_{i \in C^U} \mathbf{c}^V(t)^\top \boldsymbol{\theta}_i \end{aligned}$$

Definition 2 (Cumulative regret). The cumulative regret over T iterations of the policy $\pi \in \Pi_{C^V+U}$, is defined as $R(T) = \sum_{t=1}^T \mathbb{E}[r_{\pi^*(\mathbf{c}(t))}] - \sum_{t=1}^T \mathbb{E}[r_{\pi(\mathbf{c}(t))}]$.

Property 1 (Regret decomposition). The cumulative regret over T iterations of the policy $\pi \in \Pi_{C^V+U}$ can be rewritten as following:

$$\begin{aligned} R(T) &= \sum_{t=1}^T \mathbb{E}[r_{\pi^*(\mathbf{c}(t))}] - \sum_{t=1}^T \mathbb{E}[r_{\pi(\mathbf{c}(t))}] \\ &= \sum_{t=1}^T \left[\mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k^*(t)} - \mathbf{c}^{V+U}(t)^\top \boldsymbol{\mu}_{k(t)} \right] \\ &= \sum_{t=1}^T \left[\mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k^*(t)} - \mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k(t)} \right] \\ &\quad + \sum_{t=1}^T \left[\mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k(t)} - \mathbf{c}^{V+U}(t)^\top \boldsymbol{\mu}_{k(t)} \right] \\ &= \sum_{t=1}^T \left[\mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k^*(t)} - \mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k(t)} \right] + \\ &\quad \sum_{t=1}^T \left[\mathbb{E}[r_{U^*} | \mathbf{c}^V(t), \pi] - \mathbb{E}[r_{U(t)} | \mathbf{c}^V(t), \pi] \right] \end{aligned}$$

$$\begin{aligned} &= \sum_{t=1}^T \left[\mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k^*(t)} - \mathbf{c}^{V+U^*}(t)^\top \boldsymbol{\mu}_{k(t)} \right] + \\ &\quad \sum_{t=1}^T \left[\sum_{i \in C^{U^*}} \mathbf{c}^V(t)^\top \boldsymbol{\theta}_i - \sum_{i \in C^{U(t)}} \mathbf{c}^V(t)^\top \boldsymbol{\theta}_i \right]. \end{aligned}$$

Remark 1. CAB problem generalizes contextual bandit with restricted context problem (Bouneffouf et al., 2017). Indeed, when the subset of observed context C^V is empty, the reward of selecting the feature i is given by $\theta_{i, N+1}$, which is the coordinate $N+1$ of the vector $\boldsymbol{\theta}_i$.

Before introducing an algorithm for solving the above CAB problem, we will derive a lower bound on the expected regret of any algorithm used to solve this problem.

Theorem 1. For any policy $\pi \in \Pi_{C^V+U}$ solving Context Attentive Bandit problem (Algorithm 1) under Assumption 1.1 and 1.2, there exists probability distribution $P_r(r | \mathbf{c}, k)$, such that the lower bound of the regret accumulated by π over T iterations is :

$$\Omega \left(\sqrt{(U+V)T} + U\sqrt{VT} \right).$$

Proof. The left term of the regret (see Property 1) is lower bounded by the lower bound of linear bandits in dimension $U+V$ (Theorem 2 in [Chu et al., 2011]), while the right term is lower bounded by the lower bound of U linear bandits in dimension V . \square

4 Context Attentive Thompson Sampling (CATS)

We now propose an algorithm for solving the CAB problem, called *Context-Attentive Thompson Sampling (CATS)*, and summarize it in Algorithm 2. The basic idea of CATS is to use linear Thompson Sampling [Agrawal and Goyal, 2013] for solving the U linear bandit problems for selecting the set of additional relevant features C^{U^*} knowing $\mathbf{c}^V(t) \in C^V$, and for selecting the best arm knowing $\mathbf{c}^{V+U^*}(t) \in C^{V+U^*}$. Linear Thompson Sampling assumes a Gaussian prior for the likelihood function, which corresponds in CAB for arm k to $\tilde{\boldsymbol{\mu}}_k \sim \mathcal{N}(\mathbf{c}^{V+U}(t)^\top \tilde{\boldsymbol{\mu}}_k, \alpha^2)$, and for feature i to, $\tilde{\boldsymbol{\theta}}_i \sim \mathcal{N}(\mathbf{c}^V(t)^\top \tilde{\boldsymbol{\theta}}_i, \alpha^2)$. Then the posterior at time $t+1$ are respectively $P(\tilde{\boldsymbol{\mu}}_k | r_k(t)) \sim \mathcal{N}(\tilde{\boldsymbol{\mu}}_k, \alpha^2 A^{-1} \mathbf{1}_k)$ for arm k , and $P(\tilde{\boldsymbol{\theta}}_i | r_i(t)) \sim \mathcal{N}(\tilde{\boldsymbol{\theta}}_i, \alpha^2 B^{-1} \mathbf{1}_i)$ for feature i .

The algorithm takes the total number of features N , the number of features initially observed V , the number of additional features to observe U , the set of observed features C^V , the number of actions K , the time horizon T , the distribution parameter $\alpha > 0$ used in linear Thompson Sampling, and a function of time $\lambda(t)$, which is used for adapting the algorithm to non-stationary linear bandits.

At each iteration t the values $\mathbf{c}^V(t)$ of features in the subset C^V are observed (line 4 Algorithm 2). Then the vector parameters $\tilde{\boldsymbol{\theta}}_i$ are sampled for each feature $i \in C^V$ (lines 5-7) from the posterior distribution (line 6). Then the subset of best estimated features at time t is selected (lines 8-9). Once the feature vector $\mathbf{c}^{U(t)}$ is observed line 10, Linear Thompson

Algorithm 2 Context Attentive Thompson Sampling (CATS)

```

1: Require:  $N, V, U, K, T, C^V, \alpha > 0, \lambda(t)$ 
2: Initialize:  $\forall k \in [K], A_k := I_{N+1}, \mathbf{g}_k := \mathbf{0}_{N+1}, \hat{\boldsymbol{\mu}}_k := \mathbf{0}_{N+1}$ , and  $\forall i \in [N], B_i := I_{N+1}, \mathbf{z}_i := \mathbf{0}_{N+1}, \hat{\boldsymbol{\theta}}_i := \mathbf{0}_{N+1}$ .
3: Foreach  $t = 1, 2, \dots, T$  do
4:   observe  $\mathbf{c}^V(t)$ 
5:   Foreach context feature  $i \in C \setminus C^V$  do
6:     Sample  $\tilde{\boldsymbol{\theta}}_i$  from  $\mathcal{N}(\hat{\boldsymbol{\theta}}_i, \alpha^2 B_i^{-1})$ 
7:   End do
8:   Sort  $(\mathbf{c}^V(t)^\top \tilde{\boldsymbol{\theta}}_1, \dots, \mathbf{c}^V(t)^\top \tilde{\boldsymbol{\theta}}_N)$  in decreasing order
9:   Select  $C^U(t) := \{i \in C \setminus C^V, \mathbf{c}^V(t)^\top \tilde{\boldsymbol{\theta}}_i \geq \mathbf{c}^V(t)^\top \tilde{\boldsymbol{\theta}}_V\}$ 
10:  observe values  $\mathbf{c}^{V+U}(t)$ 
11:  Foreach arm  $k = 1, \dots, K$  do
12:    Sample  $\hat{\boldsymbol{\mu}}_k$  from  $\mathcal{N}(\hat{\boldsymbol{\mu}}_k, \alpha^2 A_k^{-1})$  distribution
13:  End do
14:  Select arm  $k(t) := \arg \max_{k \in [K]} \mathbf{c}^{V+U}(t)^\top \hat{\boldsymbol{\mu}}_k$ 
15:  Observe  $r_k(t)$ 
16:   $A_k := A_k + \mathbf{c}^{V+U}(t) \mathbf{c}^{V+U}(t)^\top, \mathbf{g}_k := \mathbf{g}_k + \mathbf{c}^{V+U}(t) r_k(t), \hat{\boldsymbol{\mu}}_k := A_k^{-1} \mathbf{g}_k$ 
17:  Foreach  $i \in C^U$ 
18:     $B_i := \lambda(t) B_i + \mathbf{c}^V(t) \mathbf{c}^V(t)^\top, \mathbf{z}_i := \mathbf{z}_i + \mathbf{c}^V(t) r_k(t), \hat{\boldsymbol{\theta}}_i := \lambda(t) B_i^{-1} \mathbf{z}_i$ 
19:  End do
20: End do

```

Sampling is applied in steps 11-15 to choose an arm. When the reward of selected arm is observed (line 15) the parameters are updated lines 16-19.

Remark 2 (Algorithmic complexity). *At each time set, Algorithm 2 sorts a set of size V and inverts $U + 1$ matrices in dimensions $N + 1$ that leads to an algorithmic complexity in $O(V \log V + (U + 1)(N + 1)^2 T)$.*

Due to assumption 2, CATS algorithm benefits from a linear algorithmic complexity, overcoming the negative result stated in [Foster *et al.*, 2016]. Before providing an upper bound of the regret of CATS in $\tilde{O}(\sqrt{T})$ (\tilde{O} hides logarithmic factors), we need an additional assumption on the noise $\eta_k(t)$.

Assumption 3 (Sub-Gaussian noise): $\forall C^{U+V} \in C$ and $\forall k \in [K]$, the noise $n_k(t) = r_k(t) |\mathbf{c}^{V+U}(t) - \mathbf{c}^{U+V}(t)^\top \hat{\boldsymbol{\mu}}_k|$ is conditionally ρ -sub-Gaussian with $\rho \geq 0$, that is for all $t \geq 1$,

$$\forall \lambda \in \mathbb{R}, \quad E[e^{\lambda n_k(t)}] \leq \exp\left(\frac{\lambda^2 \rho^2}{2}\right).$$

Lemma 1. (Theorem 2 in [Agrawal and Goyal, 2013]) *When the measurement noise $n_k(t)$ satisfies Assumption 2, $\forall t \in \{1, \dots, T\}$ $\|\mathbf{c}_k(t)\| \leq 1, \|\boldsymbol{\mu}\| \leq 1, \alpha = \rho \sqrt{9N \log \frac{T}{\delta}}$, and $\epsilon = \frac{1}{\log T}$ the regret $R(T)$ of Thompson Sampling in the Linear bandit problem with K parameters is upper bounded with a probability $1 - \delta$ by:*

$$O\left(N \sqrt{KT \log K} (\log T)^{3/2} \log \frac{1}{\delta}\right), \text{ where } 0 < \delta < 1.$$

We can now derive the following result.

Theorem 2. *When the measurement noise $\eta_k(t)$ satisfies Assumption 3, $\forall t \in \{1, \dots, T\}$ $\|\mathbf{c}_k(t)\| \leq 1, \|\boldsymbol{\mu}\| \leq 1, \alpha = \rho \sqrt{9N \log \frac{T}{\delta}}, \lambda(t) = 1$ and $\epsilon = \frac{1}{\log T}$ the regret $R(T)$ of CATS (Algorithm 2) is upper bounded with a probability $1 - \delta$ by:*

$$O\left(\left((U+V)\sqrt{KT \log K} + UV \sqrt{(N-V)T \log(N-V)}\right) (\log T)^{3/2} \log \frac{1}{\delta}\right).$$

Proof. For upper bounding the left term of the regret (see Property 1), we apply Lemma 1, and for upper bounding the right term, which is the regret of Thompson Sampling in U linear bandit problems in V dimensions with respectively $N - V, N - V - 1, \dots, N - V - U - 1$ parameters, we apply Lemma 1. \square

Theorem 2 states that the regret of CATS depends on the following two terms: the left term is the regret due to selecting a sub-optimal arm, while the right term is the regret of selecting a sub-optimal subset of features. We can see that there is still a gap between the lower bound of the Context Attentive Bandit problem and the upper bound of the proposed algorithm. The left term of the lower bound scales in $\Omega(\sqrt{(U+V)T})$, while the left term of the upper bound of CATS scales in $\tilde{O}((U+V)\sqrt{KT \log K})$, where \tilde{O} hides logarithmic factor. The right term of the lower bound scales in $\Omega(U\sqrt{VT})$, while the right term of the upper bound of CATS scales in $\tilde{O}((U+V)\sqrt{(N-V)T \log(N-V)})$. These gaps are due to the upper bound of regret of CATS, which uses Lemma 1. This suggests that the use of linear bandits based on an upper confidence balls, which scale in $\tilde{O}\sqrt{dT}$ [Abbasi-Yadkori *et al.*, 2011] (d is the dimension of contexts), could reduce this theoretical gap. As we show in the next section, we choose the Thompson Sampling approach for its better empirical performances.

5 Experiments

We compare the proposed CATS algorithm with:

(1) Random-EI: in addition to the V observed features, this algorithm selects a Random subset of features of the specified size U at each Iteration (thus, Random-EI), and then invokes the linear Thompson sampling algorithm.

(2) Random-fix: this algorithm invokes linear Thompson sampling on a subset of $U + V$ features, where the subset V is randomly selected once prior to seeing any data samples, and remains fixed.

(3) The state-of-art method for context-attentive bandits proposed in [Bouneffouf *et al.*, 2017], Thompson Sampling with Restricted Context (TSRC): TSRC solves the CBRC (contextual bandit with restricted context) problem discussed earlier: at each iteration, the algorithm decides on $U + V$ features to observe (referred to as *unobserved context*). In our setting, however, V out of N features are immediately observed at each iteration (referred to as *known context*), then

TSRC decision mechanism is used to select U additional unknown features to observe, followed by linear Thompson sampling on $U+V$ features.

(4) CALINUCB: where we replace the contextual TS in CATS with LINUCB.

(5) CATS-fix: is heuristic where we stop the features exploration after some iterations $T' = 10\%, 20\% \dots 90\%$ (we report here the an average over the best results).

Empirical evaluation of Random-fix, Random-EI, CATS-fix, CALINUCB, CATS¹ and TSRC was performed on several publicly available datasets, as well as on a proprietary corporate dialog orchestration dataset. Publicly available Covertypes and CNAE-9 were featured in the original TSRC paper and Warfarin [Sharabiani *et al.*, 2015] is a historically popular dataset for evaluating bandit methods.

To simulate the known and unknown context space, we randomly fix 10% of the context feature space of each dataset to be known at the onset and explore a subset of U unknown features. To consider the possibility of nonstationarity in the unknown context space over time, we introduce a weight decay parameter $\lambda(t)$ that reduces the effect of past examples when updating the CATS parameters. We refer to the stationary case as CATS and fix $\lambda(t) = 1$. For the nonstationary setting, we simulate nonstationarity in the unknown feature space by duplicating each dataset, randomly fixing the known context in the same manner as above, and shuffling the unknown feature set - label pairs. Then we stochastically replace events in the original dataset with their shuffled counterparts, with the probability of replacement increasing uniformly with each additional event. We refer to the nonstationary case as NCATS and use $\lambda(t)$ as defined by the GP-UCB algorithm [Srinivas *et al.*, 2009]. We compare NCATS to NCATS-fix and NCALINUCB which are the non stationary version of CATS-fix and CALINUCB. we have also compare NCATS to the Weighted TSRC (WTSRC), the nonstationary version of TSRC also developed by [Bouneffouf *et al.*, 2017]. WTSRC makes updates to its feature selection model based only on recent events, where recent events are defined by a time period, or "window" w . We choose $w = 100$ for WTSRC. We report the total average reward divided by T over 200 trials across a range of U corresponding to various percentages of N for each algorithm in Table ??.

The results in Table ?? are promising, with *our methods outperforming the state of the art in the majority of cases across both settings*. The most notable exception is found for CNAE-9 dataset, where CATS sometimes outperforms or nearly matches TSRC performance. This outcome is somewhat expected, since in the original work on TSRC [Bouneffouf *et al.*, 2017], the mean error rate of TSRC was only 0.03% lower than the error corresponding to randomly fixing a subset of unknown features to reveal for each event on CNAE-9. This observation suggests that, for this particular dataset, there may not be a subset of features which would be noticeably more predictive of the reward than the rest of the features. We also observe that the LINUCB version of

CATS has comparable performance with CATS with slight advantage to CATS. Another observation is that CATS-fix is performing better than CATS in some situations, the explanation could be that after finding the best features the algorithm do not need to explore anymore and focus on finding the best arms based on these features. Note that there may not be a subset of features which would be noticeably more predictive of the reward than the rest of the features, which is the underlying assumption of TSRC. On top of this assumption, CATS also assumes that there exist relationships between the known and unknown context features, which is likely to cause a small compounding of error.

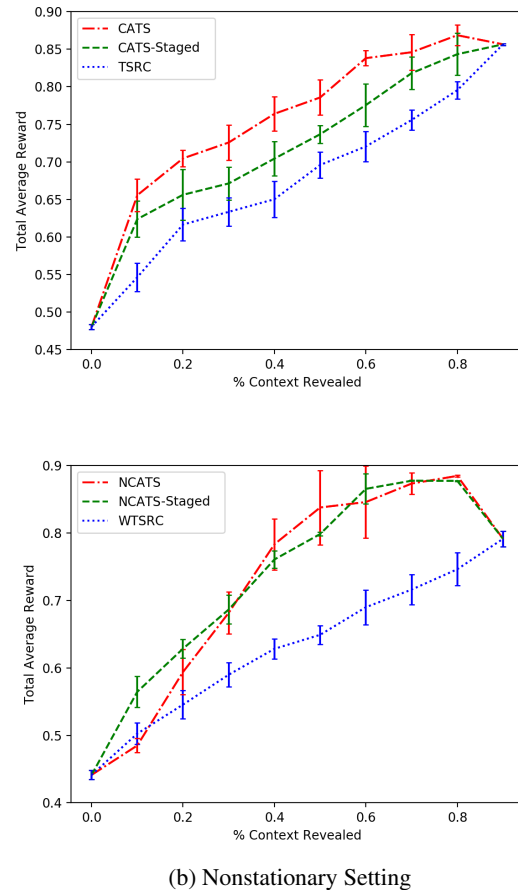


Figure 1: Total Average Reward for Covertypes

We perform a deeper analysis of the Covertypes dataset, examining multi-staged selection of the U unknown context feature sets. In CATS, the known context is used to select all U additional context feature sets at once. In a multi-staged approach, the known context grows and is used to select each of the U additional context features incrementally (one feature at a time). Maintaining $\lambda(t) = 1$, for the stationary case we denote these two cases of the CATS algorithm as CATS and CATS-Staged respectively and report their performance when 10% of the context is randomly fixed, across various U in Figure 1a.

¹Note that we have used the same exploration parameter value used in [Chapelle and Li, 2011] for TS and LINUCB type algorithms which are $TS \in \{0.25, 0.5, 1\}$ and $LINUCB \in \{0.51, 1, 2\}$

| (a) Stationary setting | | | | (b) Nonstationary setting | | | |
|------------------------|---------------------|---------------------|---------------------|---------------------------|---------------------|---------------------|---------------------|
| | <i>Warfarin</i> | | | | <i>Warfarin</i> | | |
| <i>U</i> | 20% | 40% | 60% | <i>U</i> | 20% | 40% | 60% |
| TSRC | 53.28 ± 1.08 | 57.60 ± 1.16 | 59.87 ± 0.69 | WTSRC | 55.83 ± 0.55 | 58.00 ± 0.83 | 59.85 ± 0.60 |
| CATS | 53.65 ± 1.21 | 58.55 ± 0.67 | 60.40 ± 0.74 | NCATS | 59.47 ± 2.89 | 59.34 ± 2.04 | 63.26 ± 0.75 |
| CATS-fix | 53.99 ± 1.02 | 58.67 ± 0.65 | 60.07 ± 0.54 | NCATS-fix | 59.01 ± 3.09 | 59.14 ± 2.33 | 62.42 ± 0.98 |
| CALINUCB | 52.17 ± 0.89 | 57.23 ± 0.53 | 60.29 ± 0.66 | NCLINUCB | 58.64 ± 2.77 | 58.43 ± 1.89 | 63.01 ± 0.66 |
| Random-fix | 51.05 ± 1.31 | 53.55 ± 0.97 | 55.15 ± 0.83 | Random-fix | 43.91 ± 1.17 | 47.67 ± 1.08 | 54.18 ± 1.03 |
| Random-EI | 43.65 ± 1.21 | 48.55 ± 1.67 | 50.40 ± 1.33 | Random-EI | 47.78 ± 2.11 | 52.55 ± 1.83 | 55.45 ± 1.54 |
| | <i>Covertypes</i> | | | | <i>Covertypes</i> | | |
| <i>U</i> | 20% | 40% | 60% | <i>U</i> | 20% | 40% | 60% |
| TSRC | 54.64 ± 1.87 | 63.35 ± 1.87 | 69.59 ± 1.72 | WTSRC | 50.26 ± 1.58 | 58.99 ± 1.81 | 64.91 ± 1.38 |
| CATS | 65.57 ± 2.17 | 72.58 ± 2.36 | 78.58 ± 2.35 | NCATS | 48.50 ± 1.05 | 68.17 ± 3.14 | 83.78 ± 5.51 |
| CATS-fix | 65.88 ± 2.01 | 72.67 ± 2.13 | 78.55 ± 2.25 | NCATS-fix | 49.87 ± 1.20 | 68.04 ± 3.24 | 82.98 ± 5.83 |
| CALINUCB | 61.99 ± 1.53 | 72.54 ± 1.76 | 79.69 ± 1.82 | NCLINUCB | 48.12 ± 0.99 | 68.20 ± 3.11 | 83.01 ± 5.21 |
| Random-fix | 53.11 ± 1.45 | 59.67 ± 1.07 | 64.18 ± 1.03 | Random-fix | 43.11 ± 3.05 | 49.67 ± 2.77 | 53.18 ± 2.33 |
| Random-EI | 46.15 ± 2.61 | 52.55 ± 1.81 | 55.45 ± 1.5 | Random-EI | 44.45 ± 4.44 | 46.65 ± 3.88 | 53.45 ± 3.61 |
| | <i>CNAE-9</i> | | | | <i>CNAE-9</i> | | |
| <i>U</i> | 20% | 40% | 60% | <i>U</i> | 20% | 40% | 60% |
| TSRC | 33.57 ± 2.43 | 38.62 ± 1.68 | 42.05 ± 2.14 | WTSRC | 19.91 ± 2.67 | 30.86 ± 2.92 | 36.01 ± 2.88 |
| CATS | 29.84 ± 1.82 | 39.10 ± 1.41 | 40.52 ± 1.42 | NCATS | 30.88 ± 0.96 | 34.91 ± 1.93 | 42.04 ± 1.52 |
| CATS-fix | 29.82 ± 1.70 | 39.57 ± 1.23 | 41.43 ± 1.39 | NCATS-fix | 29.92 ± 1.06 | 33.43 ± 1.83 | 40.04 ± 1.51 |
| CALINUCB | 28.53 ± 1.65 | 38.88 ± 1.35 | 39.73 ± 1.36 | NCLINUCB | 31.07 ± 0.87 | 34.61 ± 1.73 | 41.81 ± 1.62 |
| Random-fix | 33.01 ± 1.82 | 37.67 ± 1.68 | 39.18 ± 1.52 | Random-fix | 13.01 ± 3.45 | 21.77 ± 3.08 | 24.18 ± 2.43 |
| Random-EI | 32.05 ± 2.01 | 36.65 ± 1.90 | 37.47 ± 1.75 | Random-EI | 16.15 ± 2.44 | 22.55 ± 2.18 | 25.45 ± 2.15 |

 Table 1: Total average reward, $V = 10\%$

Note that when the remaining 90% of features are revealed, the CATS and TSRC methods all reduce to simple linear Thompson sampling with the full feature set. Similarly, when 0 additional feature sets are revealed, the methods all reduce to linear Thompson sampling with a sparsely represented known context. Observe that CATS consistently outperforms CATS-Staged across all U tested.

CATS-Staged likely suffers because incremental feature selection adds nonstationarity to the known context - CATS learns relationships between the known and unknown features while CATS-Staged learns relationships between them as the known context grows. Nonetheless, both methods outperform TSRC. In the nonstationary case we use the GP-UCB algorithm for $\lambda(t)$, refer to the single and multi-staged cases as NCATS and NCATS-Staged, and illustrate their performance in Figure 1b. Here we observe that NCATS and NCATS-Staged have comparable performance, and the improvement gain over baseline, in this case WTSRC, is even greater than in the stationary case.

5.1 Customer Assistant Evaluation



Next we evaluate our methods on *Customer Assistant*, a proprietary multi-skill dialog orchestration dataset. Recall that this kind of application motivates the CAB setting because there is a natural divide between the known and unknown context spaces; the query and its associated features are known at the onset and the potential responses and their associated features are only known for the domain specific agents the query is posed to. In this case, nonstationarity in the unknown context space could reflect independent updates to the underlying domain specific agents, perhaps improving or expanding their responses. As a result, similar queries, defining the known context, could elicit vastly different distributions of response features, the unknown context, over time.

The Customer Assistant orchestrates 9 domain specific agents which we arbitrarily denote as $Skill_1, \dots, Skill_9$ in the discussion that follows. In this application, example skills lie in the domains of payroll, compensation, travel, health benefits, and so on. In addition to a textual response to a user query, the skills orchestrated by *Customer Assistant* also return the following features: an *intent*, a short string descriptor that categorizes the perceived intent of the query, and a *confidence*, a real value between 0 and 1 indicating how confident a skill is that its response is relevant to the query. Skills have multiple intents associated with them. The orchestrator uses all the features associated with the query and the candidate responses from all the skills to choose which skill should carry the conversation.

The *Customer Assistant* dataset contains 28,412 events associated with a correct skill response. We encode each query by averaging 50 dimensional GloVe word embeddings [Pennington *et al.*, 2014] for each word in each query and for each skill we create a feature set consisting of its confidence and a one-hot encoding of its intent. The skill feature set size for $Skill_1, \dots, Skill_9$ are 181, 9, 4, 7, 6, 27, 110, 297, and 30 respectively.

We concatenate the query features and all of the skill features to form a 721 dimensional context feature vector for each event in this dataset. Recall that there is no need for simulation of the known and unknown contexts; in a live setting the query features are immediately calculable or known, whereas the confidence and intent necessary to build a skill's feature set are unknown until a skill is executed. Because the confidence and intent for a skill are both accessible post execution, we reveal them together. We accommodate this by slightly modifying the objective of CATS to reveal U unknown skill feature sets instead of U unknown individual features for each event. We perform a deeper analysis of the *Cus-*

tomers Assistant dataset, examining multi-staged selection of the U unknown context feature sets. Maintaining $\lambda(t) = 1$, for the stationary case the results are summarized in Figure 2a. Here both CATS-Staged and CATS methods outperform TSRC by a large margin.



(b) Nonstationary Setting

Figure 2: Total Average Reward for Customer Assistant

For the nonstationary case we simulate nonstationarity in the same manner as the publicly available datasets, except using the natural partition of the query features as the known context and the skill feature sets as the unknown context instead of simulated percentages. We use the GP-UCB algorithm for $\lambda(t)$ and illustrate the performance of NCATS and NCATS-Staged alongside WTSRC in Figure 2b. Here we observe that NCATS slightly outperforms NCATS-Staged, and both outperform the WTSRC baseline.

6 Conclusions and Future Work

We have introduced here a novel bandit problem with only partially observable context and the option of requesting a limited number of additional observations. We also propose an algorithm, designed to take an advantage of the initial partial observations in order to improve its choice of which additional features to observe, and demonstrate its advantages over the prior art, a standard context-attentive bandit with no

partial observations of the context prior to feature selection step. Our problem setting is motivated by several realistic scenarios, including medical applications as well as multi-domain dialog systems. Note that our current formulation assumes that all unobserved features have equal observation cost. However, a more practical assumption is that some features may be more costly than others; thus, in our future work, we plan to expand this notion of budget to accommodate more scenarios involving different feature costs. Other directions for future work include using non-bandit algorithms in the context feature selection stage.

7 Broader Impact

This problem has broader impacts in several domains such as voice assistants, healthcare and e-commerce.

- *Better medical diagnosis.* In a clinical setting, it is often too costly or infeasible to conduct all possible tests; therefore, given the limit on the number of tests, the doctor must decide which subset of tests will result into maximally effective treatment choice in an iterative manner. A doctor may first take a look at patient’s medical record to decide which medical test to perform, before choosing a treatment plan.
- *Better user preference modeling.* Our approach can help to develop better chatbots and automated personal assistants. For example, following a request such as, for example, “play music”, an AI-based home assistant must learn to ask several follow-up questions (from a list of possible questions) to better understand the intent of a user and to remove ambiguities: e.g., what type of music do you prefer (jazz, pop, etc)? Would you like it on hi-fi system or on TV? And so on. Another example: a support desk chatbot, in response to user’s complaint (“My Internet connection is bad”) must learn to ask a sequence of appropriate questions (from a list of possible connection issues): how far is your WIFI hotspot? Do you have a 4G subscription? These scenarios are well-handled by the framework we proposed in this paper.
- *Better recommendations.* Voice assistants and recommendation systems in general tend to lock us in our preferences, which can have deleterious effects: e.g., recommendations based only on the past history of user’s choices may reinforce certain undesirable tendencies, e.g., suggesting an online content based on a user’s with particular bias (e.g., racist, sexist, etc). On the contrary, our approach could potentially help a user to break out of this loop, by suggesting the items (e.g. news) on additional questions (additional features) which can be used to broaden user’s horizons.

References

[Abbasi-Yadkori *et al.*, 2011] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.

- [Abbasi-Yadkori *et al.*, 2012] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2012*, pages 1–9, 2012.
- [Agarwal *et al.*, 2009] Deepak Agarwal, Bee-Chung Chen, and Pradheep Elango. Explore/exploit schemes for web content optimization. In *Proceedings of the 2009 Ninth IEEE International Conference on Data Mining, ICDM '09*, pages 1–10, Washington, DC, USA, 2009. IEEE Computer Society.
- [Agrawal and Goyal, 2013] Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *ICML (3)*, pages 127–135, 2013.
- [Allesiardo *et al.*, 2014] Robin Allesiardo, Raphaël Féraud, and Djallel Bouneffouf. A neural networks committee for the contextual bandit problem. In *Neural Information Processing - 21st International Conference, ICONIP 2014, Kuching, Malaysia, November 3-6, 2014. Proceedings, Part I*, pages 374–381, 2014.
- [Auer *et al.*, 2003] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multi-armed bandit problem. *SIAM J. Comput.*, 32(1):48–77, January 2003.
- [Balakrishnan *et al.*, 2019a] Avinash Balakrishnan, Djallel Bouneffouf, Nicholas Mattei, and Francesca Rossi. Incorporating behavioral constraints in online AI systems. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 3–11. AAAI Press, 2019.
- [Balakrishnan *et al.*, 2019b] Avinash Balakrishnan, Djallel Bouneffouf, Nicholas Mattei, and Francesca Rossi. Using multi-armed bandits to learn ethical priorities for online AI systems. *IBM J. Res. Dev.*, 63(4/5):1:1–1:13, 2019.
- [Bouneffouf and Rish, 2019] Djallel Bouneffouf and Irina Rish. A survey on practical applications of multi-armed and contextual bandits. *CoRR*, abs/1904.10040, 2019.
- [Bouneffouf *et al.*, 2017] Djallel Bouneffouf, Irina Rish, Guillermo A. Cecchi, and Raphaël Féraud. Context attentive bandits: Contextual bandit with restricted context. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 1468–1475, 2017.
- [Bouneffouf *et al.*, 2019] Djallel Bouneffouf, Srinivasan Parthasarathy, Horst Samulowitz, and Martin Wistuba. Optimal exploitation of clustering and history information in multi-armed bandit. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 2016–2022. ijcai.org, 2019.
- [Carpentier and Munos, 2012] Alexandra Carpentier and Rémi Munos. Bandit theory meets compressed sensing for high dimensional stochastic linear bandit. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2012*, pages 190–198, 2012.
- [Cesa-Bianchi *et al.*, 2011] Nicolò Cesa-Bianchi, Shai Shalev-Shwartz, and Ohad Shamir. Efficient learning with partially observed attributes. *J. Mach. Learn. Res.*, 12(null):2857–2878, November 2011.
- [Chapelle and Li, 2011] Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257, 2011.
- [Chu *et al.*, 2011] Wei Chu, Lihong Li, Lev Reyzin, and Robert E. Schapire. Contextual bandits with linear payoff functions. In Geoffrey J. Gordon, David B. Dunson, and Miroslav Dudik, editors, *AISTATS*, volume 15 of *JMLR Proceedings*, pages 208–214. JMLR.org, 2011.
- [Durand and Gagné, 2014] Audrey Durand and Christian Gagné. Thompson sampling for combinatorial bandits and its application to online feature selection. In *AAAI, Workshop Sequential Decision-Making with Big Data.*, 2014.
- [Foster *et al.*, 2016] Dean Foster, Satyen Kale, and Howard Karloff. Online sparse linear regression. In *COLT*, 2016.
- [Langford and Zhang, 2008] John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in neural information processing systems*, pages 817–824, 2008.
- [Li *et al.*, 2010] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web, WWW '10*, pages 661–670, USA, 2010. ACM.
- [Oswal *et al.*, 2020] Urvashi Oswal, Aniruddha Bhargava, and Robert Nowak. Linear bandits with feature feedback. *AAAI*, 2020.
- [Pennington *et al.*, 2014] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [Qin *et al.*, 2014] Lijing Qin, Shouyuan Chen, and Xiaoyan Zhu. Contextual combinatorial bandit and its application on diversified online recommendation. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, pages 461–469. SIAM, 2014.
- [Sharabiani *et al.*, 2015] Ashkan Sharabiani, Adam Bress, Elnaz Douzali, and Houshang Darabi. Revisiting warfarin dosing using machine learning techniques. *Computational and mathematical methods in medicine*, 2015, 2015.
- [Srinivas *et al.*, 2009] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.