

---

# Nonparametric Contextual Bandits in an Unknown Metric Space

---

Nirandika Wanigasekara  
Computer Science  
National University of Singapore  
nirandiw@comp.nus.edu.sg

Christina Lee Yu  
Operations Research and Information Engineering  
Cornell University  
cleeyu@cornell.edu

## Abstract

Consider a nonparametric contextual multi-arm bandit problem where each arm  $a \in [K]$  is associated to a nonparametric reward function  $f_a : [0, 1] \rightarrow \mathbb{R}$  mapping from contexts to the expected reward. Suppose that there is a large set of arms, yet there is a simple but unknown structure amongst the arm reward functions, e.g. finite types or smooth with respect to an unknown metric space. We present a novel algorithm which learns data-driven similarities amongst the arms, in order to implement adaptive partitioning of the context-arm space for more efficient learning. We provide regret bounds along with simulations that highlight the algorithm's dependence on the local geometry of the reward functions.

## 1 Introduction

Contextual multi-arm bandits have been used to model the task of sequential decision making in which the rewards of different decisions must be learned over trial via trial-and-error. The decision maker receives reward for each of the arms (i.e. actions or options) she chooses across the time horizon  $T$ . In each trial  $t$ , the decision maker observes the context  $x_t$ , which represents the set of observable factors of the environment that could impact the performance of the action she chooses. The decision maker must select an action based on the context and all past observations. Upon choosing action  $a \in [K]$ , she observes a reward, which is assumed to be a stochastic observation of  $f_a(x)$ , the expected reward of action  $a$  at context  $x$ . In each trial, she faces the dilemma of whether to choose an action in order to learn about its performance (i.e. exploration), or to choose an action that she believes will perform well as estimated from the limited previous data (i.e. exploitation).

Consider a setting when the number of actions is very large, e.g. there is a large number of users and products on an e-commerce platform such that fully exploring the entire space of possible recommendations is costly. It is often the case that there is additional structure amongst the large space of actions, which the algorithm could exploit to learn more efficiently. In real-world applications however, this additional structure is often unknown a priori and must be learned from the data, which itself could be costly as well. It becomes important to understand the tradeoff and costs of learning relationships amongst the arm from data over the course of the contextual bandit time horizon. We consider a stochastic nonparametric contextual bandit setting in which the algorithm is not given any information a priori about the relationship between the actions. The key question is: *Can an algorithm exploit hidden structure in a nonparametric contextual bandit problem with no a priori knowledge of the underlying metric?*

**Contributions** To our knowledge, we propose the first nonparametric contextual multi-arm bandit algorithm that incorporates latent arm similarities in a setting where no a priori information about the features or metric amongst the arms is given to the algorithm. The algorithm can learn more efficiently by sharing data across similar arms, but the tradeoff between the cost of estimating arm

similarities must be carefully accounted for. Our algorithm builds upon Slivkin’s Zooming algorithm [22], adaptively partitioning the context-arm space using pairwise arm similarities estimated from the data. The adaptive partitioning allows the algorithm to naturally adapt the precision of its estimates around regions of the context-arm space that are nearly optimal, enabling the algorithm to more efficiently allocate its observations to regions of high reward.

We provide upper bounds on the regret that show the algorithm’s dependence on the local geometry of the reward functions. If we let  $f^*(x) := \max_{a \in [K]} f_a(x)$  denote the optimal reward at context  $x$ , then the regret depends on how the mass of the set  $\{(a, x) : f^*(x) - f_a(x) \in (0, \delta]\}$  scales as  $\delta$  goes to zero. This set represents the  $\delta$ -optimal region of the context-arm space except for the exactly optimal arms, i.e. the local measure of nearly optimal options centered around the optimal policy. The scaling of this set captures the notion of “gap” used in classical multi-arm bandit problems, but in the general contextual bandit setting with a large number of arms, it may be reasonable that the second optimal arm is very close in value to the optimal arm such that the gap is always very small. Instead the relevant quantity is the relative measure of arms that are  $\delta$ -optimal yet not optimal, i.e. have gap less than  $\delta$ . If the mass of such arms decreases linearly with respect to  $\delta$ , then we show that our algorithm achieves regret of  $O(\sqrt{KT})$ .

An interesting property of our algorithm is that it is fully data-dependent and thus does not depend on the chosen representation of the arm space. The arm similarities (or distances) are measured from data collected by the algorithm itself, and thus approximates a notion of distance that is defined with respect to the reward functions  $\{f_a\}_{a \in [K]}$ . The algorithm would perform the same for any permutation of the arms. In contrast, consider existing algorithms which assume a given distance metric or kernel function which the reward function is assumed to be smooth with respect to. Those algorithms are sensitive to the metric or kernel given to it, which itself could be expensive to learn or approximate from data. Suppose that nature applied a measure preserving transformation to the arm metric space such that the function is still Lipschitz but has a significantly larger Lipschitz constant. For example, consider a periodic function that repeats across the arm metric space. The performance of existing algorithms would degrade with poorer arm feature representations, whereas the algorithm we propose would remain agnostic to such manipulations.

We provide simulations that compare our algorithm to oracle variants that have special knowledge of the arms and a naive benchmark that learns over each arm separately. Initially our algorithm has a high cost due to learning the similarities, but for settings with a large number of arms and a long time horizon, the learned similarities pay off and improve the algorithm’s long run performance.

**Related Work** As there is a vast literature on multi-arm bandits, we specifically focus on literature related to the stochastic contextual bandit problem, with an emphasis on nonparametric models. In contextual bandits, in each trial the learner first observes a feature vector, refer to as “context”, associated with each arm. The optimal reward is measured with respect to the context revealed at the beginning of each trial. One approach is to directly optimize and learn over a given space of policies rather than learn the reward functions [3, 5, 12, 14]. These methods do not require strict assumptions on the reward functions but instead depend on the complexity or size of the model class.

We focus on the alternative approach of approximating reward functions, which then depend on assumptions about the structure of the reward function. A common assumption to make is that the reward function is linear with respect to the observed context vector [15, 1, 2, 13], such that the reward estimation task reduces to learning coefficient vectors of the reward functions. [2] incorporates sparsity assumptions for the high dimensional covariate setting, and [13] imposes low rank assumptions on the coefficient vectors to reduce the effective dimension.

In the linear bandit setting with  $K$  arms but only  $\Theta$  arm types for  $\Theta \ll K$ , Gentile et al proposed an adaptive clustering algorithm which maintains an undirected graph between the arms and progressively erase edges when the estimated coefficient vectors of the pair of arms is above a set threshold [6]. Two arms of the same type are assumed to have the same coefficient vector. The threshold is chosen as a function of the minimum separation condition between coefficients vectors of different types, such that eventually the graph converges to  $\Theta$  connected components corresponding to the  $\Theta$  types. Collaborative filtering bandits [16] applies the same adaptive clustering concept to the recommendation system setting where both users and item types must be learned.

In the nonparametric setting, instead of fixing a parametric model class such as linear, most work imposes smoothness conditions on the reward functions, and subsequently use nonparametric estima-

tors such as histogram binning,  $k$  nearest neighbor, or kernel methods to learn the reward functions [24, 20, 18, 19, 7]. As the contexts are observed, the estimator is applied to learn the reward of each arm separately, essentially assuming the number of arms is not too large. [7] provides an upper bound on regret of  $\tilde{O}(KT^{\frac{d+1}{d+2}})$ , where  $d$  is the dimension of the context space, and  $K$  is the number of arms.

The setting of continuum arm bandits has been introduced to approximate setting with very large action spaces. As there are infinitely many arms, it is common to impose smoothness with respect to some metric amongst the arms [17, 22, 8, 10, 9]. As the joint context-arm metric is known, these methods apply various smoothing techniques implemented via averaging datapoints with respect to a partitioning of the context-arm space, refining the smoothing parameter as more data is collected. [7] uses a  $k$  nearest neighbor estimator using the joint context-arm metric. The contextual zooming algorithm from Slivkins [22] was a key inspiration for our proposed algorithm; it uses the given context-arm metric to *adaptively* partition the context-arm product space [22]. This enables the algorithm to efficiently allocate more observations to regions of the context-arm space that are near optimal. When  $T$  is the time horizon and  $d$  is the covering dimension of the context-arm product space, the regret of the contextual zooming algorithm is bounded above by  $\tilde{O}(T^{\frac{d+1}{d+2}})$ .

For settings with large but finite number of arms, there are nonparametric models which assume different types information is known about the relationship amongst the arms. Gaussian process bandits use a known covariance matrix to fit a Gaussian process over the joint context-arm space [11]. Taxonomy MAB assumes that similarity structure amongst the arm is given in terms of a hierarchical tree rather than metric [21]. Deshmukh et al assume that the kernel matrix between pairs of arms is known, and they subsequently use kernel methods to estimate the reward functions. Cesa-Bianchi et al assumes that a graph reflecting arm similarities is given to the algorithm, and their algorithm subsequently uses the Laplacian matrix of the given graph to regularize their estimates of the reward functions [4]. Wu et al assumes an influence matrix amongst arms is known and used to share datapoints among connected arms in the estimation [23]. A limitation of these approaches is that they assume similarity information is provided to the algorithm either as a metric, kernel, or via a graph structure. In real-world applications, this similarity information is often not readily available and must be itself learned from the data.

## 2 Problem Statement

Assume that the context at each trial  $t \in [T]$  is sampled independently and uniformly over the unit interval,  $x_t \sim U(0, 1)$ , and revealed to the algorithm. Assume there are  $K$  arms, or options, that the algorithm can choose amongst at each trial  $t$ . If the algorithm chooses arm  $a_t$  at trial  $t$ , it observes and receives a reward  $\pi_t \in \mathbb{R}$  according to  $\pi_t = f_{a_t}(x_t) + \epsilon_t$ , where  $\epsilon_t \sim N(0, \sigma^2)$  is an i.i.d Gaussian noise term with mean 0 and variance  $\sigma^2$ , and  $f_a(x)$  denotes the expected reward for arm  $a$  as a function of the context  $x$ . We assume that each arm reward function  $f_a : [0, 1] \rightarrow [0, 1]$  is  $L$ -Lipschitz, i.e. for all  $x, x' \in [0, 1]^2$ ,  $|f_a(x) - f_a(x')| \leq L|x - x'|$ .

The goal of our problem setting is to maximize the total expected payoff  $\sum_{t=1}^T \pi_t$  over the time horizon  $T$ . We provide upper bounds on the expected contextual regret,

$$\mathbb{E}[R(T)] := \mathbb{E}\left[\sum_{t=1}^T (f^*(x_t) - f_{a_t}(x_t))\right] \quad \text{where} \quad f^*(x) := \max_{a \in [K]} f_a(x). \quad (1)$$

We would like to understand whether an algorithm can efficiently exploit latent structure amongst the arm reward functions if it exists. Although the number of arms may be large, they could be drawn from a smaller set of finite arm types. Alternatively the arms could be drawn from a continuum metric space such that the reward function is jointly Lipschitz over the context-arm space; however our algorithm would not have access to or knowledge of the underlying representation in the metric space.

## 3 Algorithm Intuition

We begin by describing an oracle algorithm that is given special knowledge of the relationship between the arms in the form of the context-arm metric. Assume that the arms are embedded into a metric space, and the function is Lipschitz with respect to that metric. The contextual zooming algorithm proposed by Slivkins in [22] reduces the large continuum arm set to the effective dimension

of the underlying metric space. Essentially, their model assumes that each arm is associated to some known parameter  $\theta_a \in [0, 1]$ , and that the expected joint payoff function is 1-Lipschitz continuous in the context-arm product space with respect to a known metric  $\mathcal{D}$ , such that for all context-arm pairs  $(x, a)$  and  $(x', a')$ ,  $|f(x, \theta_a) - f(x', \theta_{a'})| \leq \mathcal{D}((x, \theta_a), (x', \theta_{a'}))$ .

The key idea of Slivkin’s zooming algorithm is to use adaptive discretization to encourage the algorithm to obtain more refined estimates in the nearly optimal regions of the space while allowing coarse estimates in suboptimal regions of the context-arm space. The algorithm maintains a partition of the context-arm space consisting of “balls”, or sets, of various sizes. The algorithm estimates the reward function within a ball by averaging observed samples that lie within this ball. An upper confidence bound is obtained by accounting for the bias (proportional to the “diameter” of the ball due to Lipschitzness) and the variance due to averaging noisy observations within the ball. When a context arrives, the UCB rule is used to select a ball in the partition, and subsequently an arm in that ball. When the number of observations in a ball increases beyond the threshold such that the variance of the estimate is less than the bias, then the algorithm splits the ball into smaller balls, refining the partition locally in this region of the context-arm space.

The main intuition of the analysis is that the UCB selection rule guarantees (with high probability) that when a ball with diameter  $\Delta$  is selected, the regret incurred by selecting this ball is bounded above by order  $\Delta$ . As a result, this algorithm is able to exploit the arm similarities via the joint metric in order to aggregate samples of similar arms such that the estimates will converge more quickly. Subsequently the algorithm refines the estimates and subpartitions the space as needed for regions that are near optimal and thus require tighter estimates in order to allow the algorithm to narrow in on the optimal arm. The limitation of the previous Zooming algorithm is that it depends crucially on the given knowledge of the context-arm joint metric, which could be unknown in advance.

**Arm Similarity Estimation** In our model, we are not given any metric or features of the arm, thus the key question is whether it is still possible for an algorithm to exploit good structure amongst the arms if it exists. We propose an algorithm inspired by Slivkin’s contextual zooming algorithm, which also adaptively partitions the context-arm space with the goal to allow for coarse estimates that converge quickly initially, and subsequently selectively refine the partition and the corresponding estimates in regions of the context-arm space that are nearly optimal. The key challenge to deal with is determining how to subpartition amongst the arms when we do not know any underlying metric or feature space. Our algorithm estimates a similarity (or distance) from the collected data itself, and uses the data-dependent distances to cluster/subpartition amongst the arms. This concept is similar to clustering bandits which also learns data-driven similarities, except that the clustering bandits works assume linear reward functions and finite types, whereas our model and algorithm is more general for nonparametric functions and arms drawn from an underlying continuous space [6].

We want our algorithm to partition the context-arm product space into balls, or subsets, within which the maximum diameter is bounded, where diameter of a subset is defined as  $\text{diam}(\mathcal{S}) := \sup_{(x,a) \in \mathcal{S}} f_a(x) - \inf_{(x',a') \in \mathcal{S}} f_{a'}(x')$ . We consider balls  $\rho \subseteq [0, 1] \times [K]$  which have the form of  $[c_0(\rho), c_1(\rho)] \times \mathcal{A}(\rho)$ , where  $c_0(\rho) \in [0, 1]$  denotes the start of the context interval,  $c_1(\rho) \in [c_0(\rho), 1]$  denotes the end of the context interval, and  $\mathcal{A}(\rho) \subseteq [K]$  denotes the subset of arms. We use  $\Delta(\rho) := c_1(\rho) - c_0(\rho)$  to denote the “width” of the context interval pertaining to the ball  $\rho$ .

In order to figure out which set of arms to include in a “ball” such that the diameter is bounded, we ideally would like to measure the  $L_\infty$  distance with respect to the context interval of the ball,  $\max_{x \in [c_0(\rho), c_1(\rho)]} |f_a(x) - f_{a'}(x)|$ . As the functions are assumed to be Lipschitz with respect to the context space, a bound on the  $L_2$  distance also implies a bound on the  $L_\infty$ . Our algorithm approximates the  $L_2$  distance, defined with respect to an interval  $[u, v]$  according to

$$\mathcal{D}_u^v(a, a') := \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( f_a(z_i(u, v)) - f_{a'}(z_i(u, v)) \right)^2} \quad (2)$$

where  $z_i(u, v) = \left(1 - \frac{i}{200}\right)u + \frac{i}{200}v$ . This is a finite sum approximation to the integrated  $L_2$  distance between  $f_a$  and  $f_{a'}$  within the interval  $[u, v]$ .

Our algorithm uses the data collected for an arm in order to approximate the reward functions using a  $k$  nearest neighbor estimator, and subsequently uses the estimated reward functions to approximate  $\mathcal{D}_u^v$ . These approximate distances are then used to cluster the arms when subpartitioning. With high probability, we show that the diameter of the constructed balls is bounded by  $2L\Delta(\rho)$ . Our algorithm

collects extra samples to compute these distances, and a key part of the analysis is to understand when the improvement in the learning rate of the reward functions is sufficient enough to offset the cost of estimating arm distances.

## 4 Algorithm Statement

Let  $n_t(\rho) = \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho)$  denote the number of times  $\rho$  has been selected before trial  $t$ . Let  $\mu_t(\rho) = \frac{1}{n_t(\rho)} \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho) \pi_s$  denote the average observed reward from  $\rho$  before trial  $t$ . Define

$$UCB_t(\rho) = \mu_t(\rho) + 2L\Delta(\rho) + \sqrt{6\sigma^2 \ln(T)/n_t(\rho)}, \quad (3)$$

which gives an upper confidence bound for the maximum reward achievable by any context-arm pair in the ball  $\rho$ . The algorithm maintains two sets of balls,  $\mathcal{P}$  and  $\mathcal{P}^*$ , such that  $\mathcal{P} \cup \mathcal{P}^*$  is a partition of the context-arm space, i.e. all balls are disjoint and the union cover the entire space. We refer to balls in  $\mathcal{P}^*$  as flagged. They are given ultimate priority in the algorithm, until sufficient samples are collected to further subpartition this ball via clustering. We refer to balls in  $\mathcal{P}$  as “active”, within which priority is given to balls with higher upper confidence bound (UCB).

**Ball-Arm Selection Rule** In a given trial  $t$ , when the context  $x_t$  arrives, the algorithm identifies the flagged balls  $\rho \in \mathcal{P}^*$  which contain context  $x_t$ , i.e.  $x_t \in [c_0(\rho), c_1(\rho)]$ , and gives priority amongst them to balls with larger width  $\Delta(\rho)$ ,

$$\rho_t = \operatorname{argmax}_{\rho \in \mathcal{P}^*} \Delta(\rho) \mathbb{I}(x_t \in [c_0(\rho), c_1(\rho)]).$$

If there are no flagged balls in  $\mathcal{P}^*$  which contain  $x_t$ , then the algorithm selects an active ball  $\rho \in \mathcal{P}$  containing  $x_t$ , and gives priority to the ball with the highest upper confidence bound  $UCB_t(\rho)$ ,

$$\rho_t = \operatorname{argmax}_{\rho \in \mathcal{P}} UCB_t(\rho) \mathbb{I}(x_t \in [c_0(\rho), c_1(\rho)]). \quad (4)$$

When a ball  $\rho_t$  is chosen, the algorithm plays an arm  $a_t \in \mathcal{A}_{\rho_t}$  via a round robin ordering. The algorithm observes a noisy reward  $\pi_t$  for arm  $a_t$  and updates  $n_t(\rho)$ ,  $\mu_t(\rho)$ , and  $UCB_t(\rho)$  accordingly.

By grouping the context-arm pairs into balls, the algorithm aggregates the observed rewards within a ball to trade-off between bias and variance. For any given trial, the algorithm reduces the decision problem from selecting amongst a large number of arms to selecting amongst a smaller set of balls, which each consist of a subset of arms. Whenever the ball is subpartitioned, the width of the context interval is halved, such that balls never repeat, and are always strictly nested within a hierarchy. Furthermore, the fact that the algorithm gives priority to flagged balls with larger context widths implies that the data collected in the “flagged” phase of every ball will be uniformly distributed over context width of that ball.

**Flagging Rule** At the beginning of the algorithm, the entire context-arm space is flagged as a single large ball to be subpartitioned, i.e.  $\mathcal{P}^* = \{([0, 1] \times [K])\}$  and  $\mathcal{P} = \emptyset$ . In subsequent rounds, we flag a ball  $\rho \in \mathcal{P}$  whenever it satisfies the condition  $n_t(\rho_t) > 6\sigma^2 \ln(T)/L^2 \Delta^2(\rho_t)$ . Upon being flagged,  $\rho$  is removed from  $\mathcal{P}$  and added to  $\mathcal{P}^*$ . Let stopping time  $\tau_f(\rho)$  denote the trial  $t$  that ball  $\rho$  is flagged. Intuitively, the threshold is chosen at a point where the confidence radius, i.e. natural variation in the estimates due to the additive Gaussian observation error, is on the order of the diameter of the ball. As a result, further collecting samples does not improve the overall UCB because the diameter of the ball will dominate the expression.

**Sub-Partitioning via Clustering** Recall that flagged balls in  $\mathcal{P}^*$  are always given priority over active balls in  $\mathcal{P}$ . The observations collected in the flagged phase are used to estimate distances, or similarities between the arms for the purpose of subpartitioning the ball into smaller balls. In particular, the algorithm splits the context space  $[c_0(\rho), c_1(\rho)]$  into 64 evenly sized intervals and waits until it collects at least  $k$  samples within each of the 64 intervals for each of the arms  $a \in \mathcal{A}(\rho)$ , where  $k$  is chosen according to  $k = 5431\sigma^2 \ln(T|\mathcal{A}(\rho)|)/(L^2 \Delta^2(\rho))$ . This condition is mathematically stated as  $\prod_{a \in \mathcal{A}(\rho)} \text{SUFFDATA}(a) = 1$  where

$$\text{SUFFDATA}(a) := \prod_{i=1}^{64} \mathbb{I}\left(\sum_{s > \tau_f(\rho)} \mathbb{I}(\rho_s = \rho, a_s = a) \mathbb{I}(x_s \in [w_{i-1}, w_i]) \geq k\right),$$

for  $w_i = c_0(\rho) + i\Delta(\rho)/64$ . When this sufficient data condition is satisfied, the algorithm uses the observations collected in the flagged phase to compute pairwise arm distances approximating (2). Let  $\tau_{cl}(\rho)$  denote the trial in which the sufficient data condition is satisfied and  $\rho$  is subpartitioned.

The SUBPARTITION subroutine estimates the reward functions via a  $k$ -nearest neighbor estimator,

$$\hat{f}_a(x) = \frac{1}{k} \sum_{s=\tau_f(\rho)+1}^{\tau_{cl}(\rho)} \mathbb{I}(\rho_s = \rho, a_s = a) \mathbb{I}(x_s \in \text{k-NN}) \pi_s, \quad (5)$$

where  $x_s$  is a  $k$  nearest neighbor of  $x$  if  $\sum_{\ell=\tau_f(\rho)+1}^{\tau_{cl}(\rho)} \mathbb{I}(\rho_\ell = \rho, a_\ell = a) \mathbb{I}(|x_\ell - x| \leq |x_s - x|) \leq k$ .

Given the estimated functions  $\{\hat{f}_a\}_{a \in \mathcal{A}(\rho)}$  and a pair of arms  $a, a' \in \mathcal{A}(\rho)$ , we compute  $\hat{\mathcal{D}}_u^v(a, a')$  for intervals  $[u, v] = [c_0(\rho), (c_0(\rho) + c_1(\rho))/2]$  and  $[u, v] = [(c_0(\rho) + c_1(\rho))/2, c_1(\rho)]$  according to

$$\hat{\mathcal{D}}_u^v(a, a') := \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( \hat{f}_a(z_i(u, v)) - \hat{f}_{a'}(z_i(u, v)) \right)^2 - \frac{2\sigma^2}{k}} \quad (6)$$

where  $z_i(u, v) = (1 - \frac{i}{200})u + \frac{i}{200}v$  and the term  $\frac{2\sigma^2}{k}$  accounts for bias due to the noise.

We use the computed distances  $\hat{\mathcal{D}}_u^v(a, a')$  to subpartition  $\rho$  by clustering the arms for each half of the context interval separately. For an arbitrary ordering of the arms, we test if the next arm has distance less than  $3L(v - u)/16$  to any of the existing cluster centers. If so, we assign it to the cluster associated to the closest cluster center. Otherwise, we create a new cluster and assign this arm to be the cluster center. This results in a clustering in which all pairs of cluster centers are guaranteed to be distance  $3L(v - u)/16$  apart, and all members of a cluster must be within distance  $3L(v - u)/16$  to the cluster center. These distances are measured with respect to the data dependent estimates  $\hat{\mathcal{D}}_u^v(a, a')$ . In our analysis, we show that with high probability  $\hat{\mathcal{D}}_u^v(a, a') \approx \mathcal{D}_u^v(a, a')$ .

Once the clusters are created, then  $\rho$  is unflagged (removed from  $\mathcal{P}^*$ ) and new balls corresponding to each of the clusters for each half of the context interval are added to the active set  $\mathcal{P}$ . See the appendix for a pseudocode description of the algorithm.

## 5 Simulation

We test our algorithm on a model with 50, 100, 200 arms and a context space of  $[0, 1]$ . Each arm  $a$  corresponds to a parameter  $\theta_a$  uniformly spaced out within  $[0, 1]$ . The expected reward for arm  $a$  and context  $x$  is

$$f_a(x) := g(x, \theta_a) = 1 - |x - 4 \min_{z \in \{0, 0.5, 1\}} |\theta_a - z|.$$

This function is periodic with respect to  $\theta$ , and can be depicted as a zigzag. Our distance estimate  $\hat{\mathcal{D}}_u^v(a, a')$  approximates  $\mathcal{D}_u^v(a, a')$ , which is defined with respect  $f_a$  and  $f_{a'}$  directly and does not depend on  $\theta_a$ . Consider a measure preserving transformation that maps  $\theta_a$  to  $\phi_a = 4 \min_{z \in \{0, 0.5, 1\}} |\theta_a - z|$ , such that the reward function is equivalently described by  $\hat{f}_a(x) = 1 - |x - \phi_a|$ . An algorithm which partitions with respect to  $\mathcal{D}_u^v(a, a')$  would be agnostic to such a transformation, as opposed to an algorithm which depends on a metric defined with respect the arm's representation, which would perform worse on  $\theta_a$  than  $\phi_a$ .

We benchmark the performance of our Approx-Zooming algorithm against three variations:

- *Approx-Zooming -With-True-Reward-Function*: We give the Approx-Zooming algorithm oracle access to evaluate  $\mathcal{D}_u^v(a, a')$  at no cost, which is used to subpartition whenever a ball is flagged.
- *Approx-Zooming -With-Similarity-Metric*: We give the Approx-Zooming algorithm oracle access to evaluate  $|\theta_a - \theta_{a'}|$  at no cost, which is used to subpartition whenever a ball is flagged.
- *Approx-Zooming -With-No-Arm-Similarity*: This naive variant uses no arm similarities, estimating each arm's reward independently. The context space is adaptively partitioned via our algorithm.

We chose the model parameters that led to the highest average cumulative reward in each baseline algorithm. For all algorithms the flagging rule is set to  $n_t(\rho) \geq 4 \ln(T)/\Delta^2$ , and  $\sigma$  was set to  $1e - 2$ . For Approx-Zooming,  $k$  was set to 10. We set the number of trials  $T$  to 100,000 as all the algorithms had converged to their optimal point by then. Additional details on how the model parameters were chosen is given in Appendix F.

In figure 1, we plot the average cumulative reward over the trials, i.e.  $\frac{1}{T} \sum_{t=1}^T \pi_t$ , where  $T$  is the total number of trials and  $\pi_t \in (0, 1)$  is the reward observed in the  $t^{\text{th}}$  trial. We plot the result for the 200 arm setting with  $\sigma$  set to  $1e-2$ . As we can see, the oracle variant of the algorithm that uses the true reward function to calculate  $\mathcal{D}_u^v(a, a')$  achieves the best cumulative reward across the entire time horizon. Not surprisingly, the algorithm which learns each arm separately takes more time to converge to the optimal policy compared to all the other methods. Our Approx-Zooming algorithm has a heavy cost up front due to the clustering of the arms globally, but the algorithm improves over the time horizon as it learns the correct arm similarities. The oracle variant which uses the similarity metric  $|\theta_a - \theta_{a'}|$  performs worse than the true  $\mathcal{D}_u^v(a, a')$  variant, as it does not account for the periodic nature of the function. This supports our intuition that algorithms which depend on a given metric are sensitive to the choice of a good vs bad metric.




Figure 1: Avg. cumulative reward vs. number of trials

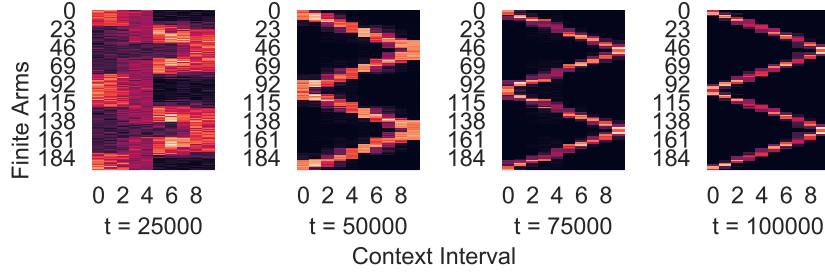


Figure 2: Approx-Zooming Selected Arm Frequency Over The Trials

In figure 2 we plot the frequencies an arm is selected in different contexts over the  $T$  trials. Each of the four plots corresponds to averaging the frequency over  $T/4$  trials across the time horizon. The x-axis refers to the context space, and the y-axis refers to the set of arms. Initially the frequency plot is very blurry, indicating that our algorithm is not necessarily playing the optimal arm but selecting arms to learn the latent arm structure. As time progresses our algorithm learns the similarities amongst arms and gradually plays the arms using the latent structure, which is depicted by the zigzag shape sharpening. Finally, in the last trials Approx-Zooming plays the optimal policy, which corresponds to the clear zigzag. In Appendix F we present similar plots for the benchmark algorithms.

Our simulations show that when the number of arms is large, it is important to use similarities amongst arms to more quickly learn the optimal policy. In addition our results highlight the fact that metric-based algorithms may be sensitive to the choice of metric, which is not a trivial task. In contrast, our approach relies on samples from the reward distribution to learn the latent structure, and is thus agnostic to any choice of metric. However, the parameter  $k$  needs to be carefully tuned for our algorithm to avoid unnecessary sampling for estimating similarities. In Appendix F we include similar plots for other parameters of the problem, in particular for smaller number of arms. We see that for 50 arms or 100 arms, the cost due to the added extra exploration may exceed the gain from learning the metric, and thus we anticipate that the benefits of learning the metric only dominates in regimes where the number of arms is large and the time horizon is sufficiently long.

## 6 Upper bound on the Regret

We present a general bound on the regret expressed as a function of a quantity relating to the local geometry of the reward function nearby the optimal policy. Let us denote  $w_i(\ell) = [(\ell - 1)2^{-i}, \ell 2^{-i}]$ ,  $\kappa(x) = f^*(x) - \max_{a \in [K]} f_a(x) \mathbb{I}(f_a(x) \neq f^*(x))$ , and

$$M_i = \sum_{\ell=1}^{2^i} \mathbb{I}(\min_{x \in w_i(\ell)} \kappa(x) \leq 20 L 2^{-i}) \sum_{a \in [K]} \mathbb{I}((f^*(2^{-i}\ell) - f_a(2^{-i}\ell)) \leq 22 L 2^{-i}).$$

**Theorem 6.1.** *The expected contextual regret of Approx-Zooming is bounded above by*

$$\mathbb{E}[R(T)] = O\left(\sigma^2 L^{-2} K \ln(TK) + \min_{i_{\max} \in \mathbb{Z}_+} (LT 2^{-i_{\max}} + \sum_{i=1}^{i_{\max}-1} \sigma^2 L^{-1} M_i 2^i \ln(TK))\right).$$

The analysis relies on showing that the instantaneous regret incurred by choosing a ball with context width  $\Delta$  is bounded above by  $O(L\Delta)$ . The first term in the regret is due to the very first initial clustering phase. The second term  $LT 2^{-i_{\max}}$  bounds the regret incurred by all balls with context width at most  $2^{-i_{\max}}$ . The terms in the summation bound the regret incurred by balls with context width equal to  $2^{-i}$ . The function  $\kappa(x)$  represents the lowest regret achieved by the second-most optimal arm, which lower bounds the suboptimality gap. In alignment with our intuition from classical MAB, when the suboptimality gap is large, the algorithm is able to more quickly converge to the optimal arm at context  $x$ . When we bound the regret incurred by all balls with context width  $2^{-i}$ , we can thus remove subintervals of the context for which  $\kappa(x)$  is large as the algorithm will have already converged to the optimal arm. This is reflected in the first indicator function within the expression  $M_i$ . Once restricted to context subintervals where the suboptimality gap is not too large, the expression  $\sum_{a \in [K]} \mathbb{I}((f^*(2^{-i}\ell) - f_a(2^{-i}\ell)) \leq 22 L 2^{-i})$  counts the number of arms for which the suboptimality gap is at most  $22 L 2^{-i}$ ; arms for which the suboptimality gap is larger will have already been deemed suboptimal. As the specific bounds on  $M_i$  depend on the model and local geometry amongst the arms, we provide bounds for two concrete examples to give more intuition.

**Finite Types** Suppose that the reward functions for the  $K$  arms,  $\{f_a\}_{a \in [K]}$  only takes  $\Theta$  different values. Essentially, this implies that there are  $\Theta$  different types of arms, but we don't know the arm types a priori. Within each type, the reward function is exactly the same. Let us define

$$\mu_\kappa(z) := \mu(\{x \in [0, 1] \text{ s.t. } \kappa(x) \leq z\})$$

where  $\mu$  is the Lebesgue measure. Then we can show that  $M_i \leq 2^i K \mu_\kappa(22 L 2^{-i})$ . The regret is bounded by the local measure function  $\mu_\kappa$ . In the finite types setting, the optimal policy corresponds to partitioning the context space  $[0, 1]$  into a set of intervals,  $\mathcal{S}^*$ , such that across each interval  $f \in \mathcal{S}^*$ , the optimal policy does not change. Let us consider the setting that  $\kappa(x)$  decreases linearly fast nearby the points where the optimal policy changes, so that for some constant  $L'$ ,  $\mu_\kappa(22 L 2^{-i}) \leq 22 |\mathcal{S}^*| L 2^{-i} / L'$ . By plugging the bound on  $M_i$  into the main theorem and choosing  $i_{\max} = \log(L' LT / 22 \sigma^2 |\mathcal{S}^*| K \ln(TK)) / 2$ , it follows that

$$\mathbb{E}[R(T)] \leq O\left(\sigma^2 L^{-2} K \ln(TK) + \sqrt{\sigma^2 |\mathcal{S}^*| L L'^{-1} T K \ln(TK)}\right) \quad (7)$$

**Lipschitz with respect to continuous arm metric space** Suppose that each arm  $a$  is associated to a latent feature  $\theta_a \in [0, 1]$ , and the expected reward function  $f_a(x) = g(x, \theta_a)$ , where  $g : [0, 1] \times [0, 1] \rightarrow [0, 1]$  is a  $L$ -Lipschitz function with respect to both the contexts and the arm latent features such that  $|g(x, \theta) - g(x', \theta')| \leq L(|x - x'| + |\theta - \theta'|)$ . If we assume that the arm latent features are uniformly spread out,  $\{\theta_a\} = \{i/K\}_{i \in [K]}$ , then

$$M_i \leq \sum_{j \in [K]} \sum_{\ell \in [2^i]} \mathbb{I}((f^*(2^{-i}\ell) - g(2^{-i}\ell, \frac{j}{K})) \leq 22 L 2^{-i}), \quad (8)$$

which is a discrete approximation to the area of the context-arm space for which the suboptimality gap is at most  $22 L 2^{-i}$ . We can visualize  $\sum_{\ell=1}^{2^i} M_i(\ell)$  by considering the contour plot of  $f^*(x) - g(x, \theta)$ , and counting how many grid points  $\{(2^{-i}\ell, j/K)\}_{\ell \in [2^i], j \in [K]}$  are lower than  $22 L 2^{-i}$ . For large  $i$  and  $K$ , this is approximately  $2^i K \mu(\{(x, \theta) : g(x, \theta) - f^*(x) \geq -22 L 2^{-i}\})$ , where  $\mu$  is the Lebesgue measure. The curve at the lowest level of the contour plot corresponds to the set  $\{(x, \theta) \text{ s.t. } g(x, \theta) -$



$f^*(x) = 0\}$ , which contains for each context the set of arm features that optimize the expected reward. The final regret depends on the local measure of the joint reward function.

As an example, if we consider the reward function  $g(x, \theta) = 1 - L|x - \theta|$  for some  $L \in (0, 1)$ , we can show that  $M_i \leq 44K$ , i.e. it is bounded by a constant with respect to  $i$ . Therefore by plugging into the main theorem and choosing  $i_{\max} = \log(20L^2T/\sigma^2K \ln(TK)) / 2$  it follows that

$$\mathbb{E}[R(T)] \leq O\left(\sigma^2 L^{-2} K \ln(TK) + \sqrt{\sigma^2 K T \ln(TK)}\right). \quad (9)$$

## 7 Discussion

**Interpreting the results.** We began this paper with the question: *Can an algorithm exploit hidden structure in a nonparametric contextual bandit problem with no a priori knowledge of the underlying metric?* The results of our simulations suggest that our proposed algorithm (with empirically tuned hyperparameters) can perform better than the corresponding algorithm that learns over each arm separately, or that uses a suboptimal metric. However, the regret bounds we present are not sufficiently strong to provably show that the algorithm outperforms learning on arms separately. The stated upper bound on regret in [7] is linear in the number of arms  $K$ , however this may be simply due to the fact that they did not optimize with respect to  $K$  in their analysis. Our regret bound is most comparable to the regret for the infinite arm setting presented in [22], and it can be recovered from their bound by imposing the discrete metric amongst the arms.

**Generalizing to higher context dimension.** For simplicity, we have stated our algorithm and analysis for the 1D context space, but the results extend to the general  $d$ -dimensional setting. The only change required algorithmically is in the subpartitioning/clustering step. Let us define  $C_d(q)$  to be the number of balls of radius  $r/q$  needed to cover a ball of radius  $r$ , which scales exponentially in the dimension  $d$ , e.g.  $q^d$ . Since we are now estimating the reward function  $f$  over a  $d$ -dimensional context space, the number of sub-regions of the context space that need to be clustered will be  $C_d(2)$ , and the number of samples needed to guarantee that the  $k$ -nearest neighbor samples are within distance  $\frac{1}{16}$  radius, will be equal to  $\tilde{O}(kC_d(32))$ . To compute  $\hat{D}$ , we will instead have a  $d$ -dimensional summation over the subset of the context space. Once  $\hat{D}$  is computed, then the clustering of arms will have the same computational cost, i.e. linear in number of arms to be clustered. The analysis can be modified to account for the  $d$ -dimensional setting, and the final regret bound will look like

$$O\left(C_d(2)C_d(32)\sigma^2 L^{-2} K \ln(TK) + \min_{i_{\max} \in \mathbb{Z}_+} \left( L 2^{-i_{\max}} + \sum_{i=1}^{i_{\max}-1} C_d(2)C_d(32)\sigma^2 L^{-1} M_i 2^i \ln(TK) \right)\right),$$

where  $M_i$  instead sums over an  $\epsilon$ -net of the context space for  $\epsilon = 2^{-i}$ , and thus we may expect  $M_i$  to grow exponentially in  $i \times d$ , depending on the distribution of the reward function and the finite arms. The growth of  $M_i$  will dominate the regret bound with respect to the dependence on the dimension  $d$ .

**Choice of metric.** Nature could apply a measure preserving transformation to the arm metric space such that the joint function has a significantly higher Lipschitz constant. This representation would incur a worse performance by the previous Zooming algorithm, indicating that the algorithm is critically dependent on the choice of representation and metric. As an example, suppose that arms are each associated to some latent parameter  $\theta \in (0, 2\pi)$ , and the reward function associated to an arm  $a$  is  $f_a(x) = x + \sin(L\theta_a)$ . The Lipschitz constant with respect to  $\theta$  is  $L$ . By applying a change of variables from  $\theta$  to  $t(\theta) = L\theta \bmod 2\pi$ , the associated reward function in terms of the representation  $t_a = t(\theta_a)$  would be  $f_a(x) = x + \sin(t_a)$ , which only has Lipschitz constant 1 with respect to the reparameterization  $t$ . This is only a simple example amongst many that illustrate the importance of the choice of metric for learning. In contrast, our algorithm estimates similarity amongst the arms directly from data collected from the reward functions, which essentially estimates distance in the function space; as a result our algorithm is invariant to any specific covariate representation.

**Future Work.** The current results are stated only for Lipschitz reward functions, where the tuning parameters depend on the Lipschitz constant. It may be interesting to generalize the algorithm to Holder continuous reward functions, and consider how to adapt the algorithm to the smoothness parameters if unknown. It would also be interesting to explore the connections to Gaussian process bandits. One would need to specify the covariance matrix amongst arms, and it may be possible to consider empirically estimating the covariance matrix in the process of learning.

## Acknowledgement

This work is supported by the National University of Singapore and A\*STAR - SERC PSF Grant 1521200084. We thank Professor David S. Rosenblum for his support of this project through insightful discussions and feedback.

## References

- [1] S. Agrawal and N. Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- [2] H. Bastani and M. Bayati. Online decision-making with high-dimensional covariates. *Available at SSRN 2661896*, 2015.
- [3] A. Beygelzimer, J. Langford, L. Li, L. Reyzin, and R. E. Schapire. Contextual bandit algorithms with supervised learning guarantees. *arXiv preprint arXiv:1002.4058*, 2010.
- [4] A. A. Deshmukh, U. Dogan, and C. Scott. Multi-task learning for contextual bandits. In *Advances in Neural Information Processing Systems*, pages 4848–4856, 2017.
- [5] M. Dudik, D. Hsu, S. Kale, N. Karampatziakis, J. Langford, L. Reyzin, and T. Zhang. Efficient optimal learning for contextual bandits. *arXiv preprint arXiv:1106.2369*, 2011.
- [6] C. Gentile, S. Li, and G. Zappella. Online clustering of bandits. In *International Conference on Machine Learning*, pages 757–765, 2014.
- [7] M. Y. Guan and H. Jiang. Nonparametric stochastic contextual bandits. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [8] E. Hazan and N. Megiddo. Online learning with prior knowledge. In *International Conference on Computational Learning Theory*, pages 499–513. Springer, 2007.
- [9] R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 681–690. ACM, 2008.
- [10] R. D. Kleinberg. Nearly tight bounds for the continuum-armed bandit problem. In *Advances in Neural Information Processing Systems*, pages 697–704, 2005.
- [11] A. Krause and C. S. Ong. Contextual gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, pages 2447–2455, 2011.
- [12] A. Krishnamurthy, J. Langford, A. Slivkins, and C. Zhang. Contextual bandits with continuous actions: Smoothing, zooming, and adapting. *arXiv preprint arXiv:1902.01520*, 2019.
- [13] S. Lale, K. Azizzadenesheli, A. Anandkumar, and B. Hassibi. Stochastic linear bandits with hidden low rank structure. *CoRR*, abs/1901.09490, 2019. URL <http://arxiv.org/abs/1901.09490>.
- [14] J. Langford and T. Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in Neural Information Processing Systems 20*, pages 817–824. 2008.
- [15] L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670. ACM, 2010.
- [16] S. Li, A. Karatzoglou, and C. Gentile. Collaborative filtering bandits. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 539–548. ACM, 2016.
- [17] T. Lu, D. Pál, and M. Pál. Showing relevant ads via context multi-armed bandits. Technical report, Tech. rep, 2009.
- [18] V. Perchet, P. Rigollet, et al. The multi-armed bandit problem with covariates. *The Annals of Statistics*, 41(2):693–721, 2013.

- [19] W. Qian and Y. Yang. Kernel estimation and model combination in a bandit problem with covariates. *The Journal of Machine Learning Research*, 17(1):5181–5217, 2016.
- [20] P. Rigollet and A. Zeevi. Nonparametric bandits with covariates. *arXiv preprint arXiv:1003.1630*, 2010.
- [21] A. Slivkins. Multi-armed bandits on implicit metric spaces. In *Advances in Neural Information Processing Systems*, pages 1602–1610, 2011.
- [22] A. Slivkins. Contextual bandits with similarity information. *The Journal of Machine Learning Research*, 15(1):2533–2568, 2014.
- [23] Q. Wu, H. Wang, Q. Gu, and H. Wang. Contextual bandits in a collaborative environment. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 529–538. ACM, 2016.
- [24] Y. Yang, D. Zhu, et al. Randomized allocation with nonparametric estimation for a multi-armed bandit problem with covariates. *The Annals of Statistics*, 30(1):100–121, 2002.

## A Algorithm Notation and Pseudocode

The algorithm is presented in full in the main body of the paper, but for reference we have also included pseudocode below.

---

### Algorithm 1 Approx-Zooming Algorithm

---

**Require:** context space  $[0, 1]$ , arm space  $[K]$ , Lipschitz constant  $L$  and time horizon  $T$

- 1: **function** APPROX-ZOOMING ( $K, L, T$ )
- 2:   Parameter  $k = \frac{5431\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{L^2\Delta^2(\rho)}$
- 3:   Initialize  $\mathcal{P}^* = \{([0, 1] \times [K])\}$ ,  $\mathcal{P} = \emptyset$
- 4:   **for**  $t \in [T]$  **do**
- 5:     Context  $x_t$  “arrives”
- 6:     **if**  $x_t \in [c_0(\rho), c_1(\rho)]$  for any  $\rho \in \mathcal{P}^*$  **then** ▷ Check for any flagged balls
- 7:        $\rho_t = \operatorname{argmax}_{\rho \in \mathcal{P}^*} \Delta(\rho) \mathbb{I}(x_t \in [c_0(\rho), c_1(\rho)])$
- 8:       Select  $a_t = \min\{a \in \mathcal{A}(\rho) \text{ s.t. } \text{SUFFDATA}(a) = 0\}$
- 9:       Play arm  $a_t$  and observe payoff  $\pi_t = f_{a_t}(x_t) + \epsilon_t$
- 10:      **if**  $\prod_{a \in \mathcal{A}(\rho)} \text{SUFFDATA}(a) == 1$ , **then**
- 11:       new partitions = SUBPARTITION( $\rho$ )
- 12:        $\mathcal{P} = \mathcal{P} \cup \text{new partitions}$
- 13:        $\mathcal{P}^* = \mathcal{P}^* \setminus \rho$
- 14:     **else** ▷ If no relevant flagged balls, then use UCB selection rule
- 15:        $\rho_t = \operatorname{argmax}_{\rho \in \mathcal{P}} \text{UCB}_t(\rho) \mathbb{I}(x_t \in [c_0(\rho), c_1(\rho)])$
- 16:       Play any arm  $a_t \in \mathcal{A}(\rho_t)$  and observe payoff  $\pi_t = f_{a_t}(x_t) + \epsilon_t$
- 17:       **if**  $n_t(\rho_t) > \frac{6\sigma^2 \ln(T)}{L^2\Delta^2(\rho_t)}$  **then**
- 18:         $\mathcal{P} = \mathcal{P} \setminus \{\rho_t\}$ ,  $\mathcal{P}^* = \mathcal{P}^* \cup \{\rho_t\}$  ▷ Remove from active set, flag for subpartition
- 19:         $\tau_f(\rho) = t$
- 20:
- 21: **function** SUBPARTITION( $\rho$ )
- 22:   Initialize centers  $\mathcal{C}_0 = \emptyset$ ,  $\mathcal{C}_1 = \emptyset$
- 23:   **for**  $\ell \in \{0, 1\}$  **do**
- 24:      $u = c_0(\rho) + \frac{\Delta(\rho)\ell}{2}$ ,  $v = c_0(\rho) + \frac{\Delta(\rho)(1+\ell)}{2}$
- 25:     **for**  $a \in \mathcal{A}_\rho$  in arbitrary order **do**
- 26:       Compute  $\hat{D}_u^v(a, y)$  for all  $y \in \mathcal{C}_\ell$
- 27:       **if**  $\min_{y \in \mathcal{C}_\ell} \hat{D}_u^v(a, y) > \frac{3}{16}L(v - u)$  **then**
- 28:        Add  $a$  as a new center,  $\mathcal{C}_\ell = \mathcal{C}_\ell \cup \{a\}$ ,  $\mathcal{S}_\ell(a) = \{a\}$
- 29:       **else**
- 30:         $y = \operatorname{argmin}_{y \in \mathcal{C}_\ell} \hat{D}_\rho(a, y)$
- 31:         $\mathcal{S}_\ell(y) = \mathcal{S}_\ell(y) \cup \{a\}$
- 32:   **return**  $\left\{ \left( [c_0(\rho), \frac{c_0(\rho)+c_1(\rho)}{2}] \times \mathcal{S}_0(y) \right) \right\}_{y \in \mathcal{C}_0} \cup \left\{ \left( [\frac{c_0(\rho)+c_1(\rho)}{2}, c_1(\rho)] \times \mathcal{S}_1(y) \right) \right\}_{y \in \mathcal{C}_1}$

---

Recall the notation and definitions introduced:

- $\text{diam}(\mathcal{S}) := \sup_{(x,a) \in \mathcal{S}} f_a(x) - \inf_{(x',a') \in \mathcal{S}} f_{a'}(x')$
- $[c_0(\rho), c_1(\rho)]$  denotes the context interval of ball  $\rho$ .
- $\mathcal{A}(\rho)$  denotes the set of arms in ball  $\rho$ .
- $\Delta(\rho) := c_1(\rho) - c_0(\rho)$  denotes the context width of ball  $\rho$ .
- $n_t(\rho) := \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho)$  denotes the number of trials ball  $\rho$  has been chosen by the algorithm before trial  $t$ .
- $\mu_t(\rho) := \frac{1}{n_t(\rho)} \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho) \pi_s$  denotes the average observed reward from this ball before trial  $t$ .
- $\text{UCB}_t(\rho) := \mu_t(\rho) + 2L\Delta(\rho) + \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}}$  is an upper confidence bound for the maximum reward achievable by any context-arm pair in the ball  $\rho$ .

- $\tau_f(\rho)$  denotes the trial that ball  $\rho$  is flagged.
- $\tau_{cl}(\rho)$  denotes the trial in which the SUBPARTITION subroutine is called.
- $\text{SUFFDATA}(a) := \prod_{i=1}^{64} \mathbb{I} \left( \sum_{s > \tau_f(\rho)} \mathbb{I}(\rho_s = \rho, a_s = a) \mathbb{I}(x_s \in [w_{i-1}, w_i]) \geq k \right)$  for  $w_i = c_0(\rho) + \frac{i\Delta(\rho)}{64}$  and  $k = 5431\sigma^2 \ln(T|\mathcal{A}(\rho)|)/(L^2\Delta^2(\rho))$ .
- The reward function for an arm  $a$  is estimated via a  $k$ -NN estimator

$$\hat{f}_a(x) = \frac{1}{k} \sum_{s=\tau_f(\rho)+1}^{\tau_{cl}(\rho)} \mathbb{I}(\rho_s = \rho, a_s = a) \mathbb{I}(x_s \in \mathbf{k}\text{-NN}) \pi_s$$

where  $x_s$  is a  $k$  nearest neighbor datapoint for computing  $\hat{f}_a(x)$  if

$$\mathbb{I}(x_s \in \mathbf{k}\text{-NN}) = \sum_{\ell=\tau_f(\rho)+1}^{\tau_{cl}(\rho)} \mathbb{I}(\rho_\ell = \rho, a_\ell = a) \mathbb{I}(|x_\ell - x| \leq |x_s - x|) \leq k.$$

- The distance between arms  $a$  and  $a'$  for interval  $[u, v]$  is estimated according to

$$\hat{\mathcal{D}}_u^v(a, a') := \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( \hat{f}_a(z_i(u, v)) - \hat{f}_{a'}(z_i(u, v)) \right)^2} - \frac{2\sigma^2}{k}$$

where  $z_i(u, v) = \left(1 - \frac{i}{200}\right)r + \frac{i}{200}s$ .

## B Proof Sketch

Our algorithm and analysis take after the Zooming algorithm [9, 22]. However, the major difference is that their model assumes the metric is directly known in advance, but our algorithm must learn the metric. In particular, each trial we subpartition a set into finer clusters, we pay extra cost to collect samples to estimate distances used for determining the subsequent clusters.

**Lemma B.1.** *With probability at least  $1 - 2T^{-1}$ , over the entire course of the algorithm for all  $t \in [T]$  and  $\rho \in \mathcal{P}$ ,*

$$\mu_t(\rho) \in \left[ \min_{(x,a) \in \rho} f_a(x) - \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}}, \max_{(x,a) \in \rho} f_a(x) + \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}} \right].$$

**Lemma B.2.** *With probability at least  $1 - 4T^{-1}$ , over the entire course of the algorithm for all trials that  $\hat{\mathcal{D}}_u^v(a, y)$  is evaluated,*

$$\left\{ |\hat{\mathcal{D}}_u^v(a, y) - \mathcal{D}_u^v(a, y)| \leq \frac{1}{8}L(v - u) \right\}.$$

We refer to the set of conditions in Lemmas B.1 and B.1 as the “good event” and denote them with  $\mathcal{G}$ .

Argue that the resulting clustering produced by SUBPARTITION satisfies that all pairs in the same cluster must be “close”.

**Lemma B.3.** *Conditioned on the “good events”  $\mathcal{G}$ , for any  $\rho \in \mathcal{P}$  at any point of the algorithm,  $\text{diam}(\rho) \leq 2L\Delta(\rho)$ .*

Let us define

$$\text{gap}(\rho) = \min_{(x,a) \in \rho} (f^*(x) - f_a(x)).$$

Conditioned on good events above, for any  $\rho \in \mathcal{P}_t$ , we upper bound the total number of trials this set can be chosen by the algorithm until either it is completely dominated, or it is flagged to be subpartitioned.

**Lemma B.4.** *Conditioned on the “good events”  $\mathcal{G}$ , if ball  $\rho$  is chosen by the algorithm at trial  $t$  via the UCB rule, then*

$$n_t(\rho) \leq \min \left\{ \frac{24\sigma^2 \ln(T)}{(\text{gap}(\rho) - 4L\Delta(\rho_t))^2}, \frac{6\sigma^2 \ln(T)}{L^2\Delta^2(\rho)} + 1 \right\} \quad (10)$$

Therefore the max number of trials ball  $\rho$  is chosen via the UCB rule is bounded above by,

$$\sum_{t=1}^{\tau_f(\rho)} \mathbb{I}(\rho_t = \rho) \leq \min \left\{ \frac{24\sigma^2 \ln(T)}{(\text{gap}(\rho) - 4L\Delta(\rho_t))^2} + 1, \frac{6\sigma^2 \ln(T)}{L^2\Delta^2(\rho)} + 2 \right\} \quad (11)$$

**Lemma B.5.** Conditioned on the “good events”  $\mathcal{G}$ , for any  $\rho \in \mathcal{P}$  such that  $\tau_f(\rho) \leq T$ ,

$$\max_{(x,a) \in \rho} (f^*(x) - f_a(x)) \leq 10L\Delta(\rho).$$

For any  $\rho \in \mathcal{P}$  that was created due to a parent in  $\mathcal{P}$  being flagged (i.e.  $\Delta(\rho) \leq \frac{1}{4}$ )

$$\max_{(x,a) \in \rho} (f^*(x) - f_a(x)) \leq 20L\Delta(\rho).$$

Note that since we assumed the reward functions are bounded in  $[0, 1]$ , the regret is always bounded by 1, thus the above upper bound is only nontrivial for  $\Delta(\rho) < 1/20L$ .

**Lemma B.6.** For any  $\rho \in \mathcal{P}$  such that  $\tau_f(\rho) \leq T$ ,

$$\mathbb{E} \left[ \sum_{t=\tau_f(\rho)+1}^T \mathbb{I}(\rho_t = \rho) \mid \tau_f(\rho) \leq T \right] \leq \frac{304 \cdot 5431 \cdot \sigma^2 |\mathcal{A}(\rho)| \ln(T|\mathcal{A}(\rho)|)}{L^2 \Delta^2(\rho)} \quad (12)$$

## C Proof of Lemmas

*Proof of Lemma B.3.* Recall from the algorithm that every ball  $\rho \in \mathcal{P}$  is constructed as a result of the SUBPARTITION routine, and is associate to a cluster with a corresponding center arm  $y \in \mathcal{A}(\rho)$ . Let us denote  $r = c_0(\rho)$  and  $s = c_1(\rho)$ , which are the endpoints used to compute distances within the SUBPARTITION subroutine. It follows that  $\Delta(\rho) = s - r$ . The algorithm enforced that for any  $a, a' \in \mathcal{A}(\rho) \times \mathcal{A}(\rho)$ , it must be that  $\hat{\mathcal{D}}_r^s(a, y) \leq \frac{3}{16}L(v - u)$  and  $\hat{\mathcal{D}}_r^s(a', y) \leq \frac{3}{16}L(v - u)$ .

Conditioned on the good events  $\mathcal{G}$ ,

$$|\hat{\mathcal{D}}_u^v(a, y) - \mathcal{D}_u^v(a, y)| \leq \frac{1}{8}L(v - u) \text{ and } |\hat{\mathcal{D}}_u^v(a', y) - \mathcal{D}_u^v(a', y)| \leq \frac{1}{8}L(v - u).$$

We can verify that  $\mathcal{D}_s^r(\cdot)$  is indeed a proper metric amongst the arms as it is simply a normalized L2 norm of the difference of associated vectors. Therefore by triangle inequality,

$$\mathcal{D}_u^v(a, a') \leq \mathcal{D}_u^v(a, y) + \mathcal{D}_u^v(a', y) \quad (13)$$

$$\leq \frac{5}{16}L(v - u) + \frac{5}{16}L(v - u) = \frac{5}{8}L(v - u). \quad (14)$$

Recall that

$$\mathcal{D}_u^v(a, a') = \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( f_a(z_i(u, v)) - f_{a'}(z_i(u, v)) \right)^2} \text{ for } z_i(u, v) = \left(1 - \frac{i}{200}\right)r + \frac{i}{200}s$$

As  $[u, v]$  is a compact set, it must achieve its supremum, and we let  $x^*$  denote the point that achieves supremum, i.e.

$$|f_a(x^*) - f_{a'}(x^*)| = \sup_{x \in [u, v]} |f_a(x) - f_{a'}(x)|.$$

As  $f_a$  and  $f_{a'}$  are  $L$ -Lipschitz, the absolute value of their difference must be  $2L$ -Lipschitz.

For  $z_i(u, v)$  as defined above, by Lipschitzness,

$$|f_a(z_i(u, v)) - f_{a'}(z_i(u, v))| \geq \max(|f_a(x^*) - f_{a'}(x^*)| - 2L|z_i(u, v) - x^*|, 0)$$

Therefore,

$$\begin{aligned} (\mathcal{D}_u^v(a, a'))^2 &= \frac{1}{200} \sum_{i \in [200]} \left( f_a(z_i(u, v)) - f_{a'}(z_i(u, v)) \right)^2 \\ &\geq \frac{1}{200} \sum_{i \in [200]} \left( \max(|f_a(x^*) - f_{a'}(x^*)| - 2L|z_i(u, v) - x^*|, 0) \right)^2 \\ &\geq \min_{x \in [u, v]} \frac{1}{200} \sum_{i \in [200]} \left( \max(|f_a(x^*) - f_{a'}(x^*)| - 2L|z_i(u, v) - x|, 0) \right)^2 \end{aligned}$$

Argue that this is lower bounded by choosing  $x = s$  (include a picture?) so that

$$|z_i(u, v) - x| = \left(1 - \frac{i}{200}\right)(v - u)$$

and by rearranging the indices we get

$$(\mathcal{D}_u^v(a, a'))^2 \geq \frac{1}{200} \sum_{i=0}^{200} \left( \max(|f_a(x^*) - f_{a'}(x^*)| - \frac{2L(v-u)i}{200}, 0) \right)^2$$

Let us define

$$i_{\max} := \min \left( \left\lfloor \frac{200|f_a(x^*) - f_{a'}(x^*)|}{2L(v-u)} \right\rfloor, 200 \right)$$

then we can lower bound  $(\mathcal{D}_u^v(a, a'))^2$  by

$$\begin{aligned} & \frac{1}{200} \sum_{i=0}^{i_{\max}} \left( |f_a(x^*) - f_{a'}(x^*)| - \frac{2L(v-u)i}{200} \right)^2 \\ &= \frac{i_{\max}+1}{200} |f_a(x^*) - f_{a'}(x^*)|^2 - \frac{2L|f_a(x^*) - f_{a'}(x^*)|(v-u)}{200^2} \sum_{i=0}^{i_{\max}} i + \frac{4L^2(v-u)^2}{200^3} \sum_{i=0}^{i_{\max}} i^2 \end{aligned}$$

Recall that

$$\sum_{i=0}^{i_{\max}} i = \frac{i_{\max}(i_{\max}+1)}{2} \leq \frac{(i_{\max}+1)^2}{2}$$

and

$$\sum_{i=0}^{i_{\max}} i^2 = \frac{i_{\max}(i_{\max}+1)(2i_{\max}+1)}{6} \geq \frac{i_{\max}^3}{3}$$

Therefore, if  $200 < \left\lfloor \frac{200|f_a(x^*) - f_{a'}(x^*)|}{2L(v-u)} \right\rfloor$ ,

$$(\mathcal{D}_u^v(a, a'))^2 \tag{15}$$

$$\geq \frac{200+1}{200} |f_a(x^*) - f_{a'}(x^*)|^2 - \frac{2L|f_a(x^*) - f_{a'}(x^*)|(v-u)(200+1)^2}{2 \cdot 200^2} + \frac{4L^2(v-u)^2}{3} \tag{16}$$

$$\geq |f_a(x^*) - f_{a'}(x^*)|^2 - \left(\frac{4}{3}\right)^2 L |f_a(x^*) - f_{a'}(x^*)|(v-u) + \frac{4}{3} L^2 (v-u)^2 \tag{17}$$

$$= \left( |f_a(x^*) - f_{a'}(x^*)| - \frac{8L(v-u)}{9} \right)^2 + \frac{44}{81} L^2 (v-u)^2 \tag{18}$$

We have shown in (14) that conditioned on the good event  $\mathcal{G}$ ,  $\mathcal{D}_u^v(a, a') \leq \frac{5}{8} L(v-u)$ , which would violate (18), as  $(\frac{5}{8})^2 \leq \frac{44}{81}$ .

Thus, conditioned on the good event  $\mathcal{G}$ , for a pair of arms  $a, a' \in \mathcal{A}(\rho) \times \mathcal{A}(\rho)$ , it must be that  $200 > \left\lfloor \frac{200|f_a(x^*) - f_{a'}(x^*)|}{2L(v-u)} \right\rfloor$ , such that

$$(\mathcal{D}_u^v(a, a'))^2 \tag{19}$$

$$\geq \frac{1}{200} \left( \frac{|f_a(x^*) - f_{a'}(x^*)|^2 200}{2L(v-u)} \right) |f_a(x^*) - f_{a'}(x^*)| \tag{20}$$

$$- \frac{L|f_a(x^*) - f_{a'}(x^*)|(v-u)}{200^2} \left( \frac{200|f_a(x^*) - f_{a'}(x^*)|}{2L(v-u)} + 1 \right)^2 \tag{21}$$

$$+ \frac{4L^2(v-u)^2}{3 \cdot 200^3} \left( \frac{200|f_a(x^*) - f_{a'}(x^*)|}{2L(v-u)} - 1 \right)^3 \tag{22}$$

$$= \frac{|f_a(x^*) - f_{a'}(x^*)|^3}{L(v-u)} \left( \frac{1}{2} - \frac{1}{4} \left( 1 + \frac{2L(v-u)}{200|f_a(x^*) - f_{a'}(x^*)|} \right) \right)^2 + \frac{1}{6} \left( 1 - \frac{2L(v-u)}{200|f_a(x^*) - f_{a'}(x^*)|} \right)^3 \tag{23}$$

Suppose that  $|f_a(x^*) - f_{a'}(x^*)| > L(v - u)$ . We would arise at a contradiction because plugging this bound into (23) would result in

$$(\mathcal{D}_u^v(a, a'))^2 > L^2(v - u)^2 \left( \frac{1}{2} - \frac{1}{4}(1.01)^2 + \frac{1}{6}(0.99)^3 \right) > \left( \frac{5}{8}L(v - u) \right)^2.$$

Therefore  $\mathcal{D}_u^v(a, a') \leq \frac{5}{8}L(v - u)$  must imply that

$$|f_a(x^*) - f_{a'}(x^*)| := \sup_{x \in [c_0(\rho), c_1(\rho)]} |f_a(x) - f_{a'}(x)| \leq L\Delta(\rho).$$

Let us denote  $(\bar{x}, \bar{a}) = \operatorname{argmax}_{(x,a) \in \rho} f_a(x)$  and  $(\underline{x}, \underline{a}) = \operatorname{argmin}_{(x,a) \in \rho} f_a(x)$ . Then

$$\operatorname{diam}(\rho) = f_{\bar{a}}(\bar{x}) - f_{\underline{a}}(\underline{x}) \tag{24}$$

$$\leq |f_{\bar{a}}(\bar{x}) - f_{\bar{a}}(\underline{x})| + |f_{\bar{a}}(\underline{x}) - f_{\underline{a}}(\underline{x})| \tag{25}$$

$$\leq L\Delta(\rho) + \max_{a, a' \in \mathcal{A}(\rho)^2} \sup_{x \in [c_0(\rho), c_1(\rho)]} |f_a(x) - f_{a'}(x)| \tag{26}$$

$$\leq 2L\Delta(\rho) \tag{27}$$

□

*Proof of Lemma B.4.* Recall that a ball is flagged (after being played) the first trial that

$$n_t(\rho) > \frac{6\sigma^2 \ln(T)}{L^2 \Delta^2(\rho)},$$

and subsequently it is removed from  $\mathcal{P}$  and no longer active, so the flagging condition is triggered exactly when

$$n_t(\rho) = \left\lfloor \frac{6\sigma^2 \ln(T)}{L^2 \Delta^2(\rho)} + 1 \right\rfloor.$$

Since the ball is played one last trial, the total number of trials the ball is played via the UCB rule over time horizon  $T$  will be

$$\sum_{t=1}^{\tau_f(\rho)} \mathbb{I}(\rho_t = \rho) = \left\lfloor \frac{6\sigma^2 \ln(T)}{L^2 \Delta^2(\rho)} + 1 \right\rfloor + 1 \leq \frac{6\sigma^2 \ln(T)}{L^2 \Delta^2(\rho)} + 2.$$

The second terms in the upper bound of Lemma B.4 are derived by considering when the ball must be “flagged”.

Next, we consider when the ball must be so suboptimal such that it is no longer chosen by the UCB rule. Conditioned on the “good events”  $\mathcal{G}$ , for all  $\rho \in \mathcal{P}$  and  $t \in [\tau_f(\rho)]$ , i.e. the ball is not yet flagged,

$$\mu_t(\rho) \in \left[ \min_{(x,a) \in \rho} f_a(x) - \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}}, \max_{(x,a) \in \rho} f_a(x) + \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}} \right]$$

such that

$$UCB_t(\rho) \in \left[ \min_{(x,a) \in \rho} f_a(x) + 2L\Delta(\rho), \max_{(x,a) \in \rho} f_a(x) + 2L\Delta(\rho) + 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}} \right].$$

Thus

$$UCB_t(\rho) \geq \min_{(x,a) \in \rho} f_a(x) + 2L\Delta(\rho) = \max_{(x,a) \in \rho} f_a(x) - \operatorname{diam}(\rho) + 2L\Delta(\rho)$$

and

$$UCB_t(\rho) \leq \max_{(x,a) \in \rho} f_a(x) + 2L\Delta(\rho) + 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}} = \min_{(x,a) \in \rho} f_a(x) + \operatorname{diam}(\rho) + 2L\Delta(\rho) + 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}}.$$

Suppose at trial  $t$ , context  $x_t$  arrives and ball  $\rho_t$  is chosen by the algorithm via the UCB rule. Let us denote  $a^* = \operatorname{argmax}_{a \in [K]} f_a(x_t)$ , and let  $\rho^*$  denote the ball which contains  $a^*$ . Then

$$UCB_t(\rho^*) \geq f^*(x_t) - \operatorname{diam}(\rho^*) + 2L\Delta(\rho^*).$$



By Lemma B.3, conditioned on the good events  $\mathcal{G}$ ,  $\text{diam}(\rho^*) \leq 2L\Delta(\rho^*)$ , thus

$$UCB_t(\rho^*) \geq f^*(x_t).$$

By conditioning on the good events  $\mathcal{G}$ , and by Lemma B.3, it also must hold that

$$UCB_t(\rho_t) \leq \min_{(x,a) \in \rho_t} f_a(x) + \text{diam}(\rho_t) + 2L\Delta(\rho_t) + 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho_t)}} \quad (28)$$

$$\leq \min_{(x,a) \in \rho_t} f_a(x) + 4L\Delta(\rho_t) + 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho_t)}}. \quad (29)$$

Thus,

$$\min_{a \in \mathcal{A}(\rho_t)} f_a(x_t) \geq UCB_t(\rho_t) - 4L\Delta(\rho_t) - 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho_t)}}.$$

If  $\rho_t$  was chosen by the algorithm via the UCB rule, it therefore must imply that

$$UCB_t(\rho) \geq UCB_t(\rho^*)$$

and thus

$$f^*(x_t) - \min_{a \in \mathcal{A}(\rho_t)} f_a(x_t) \leq 4L\Delta(\rho_t) + 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho_t)}}.$$

This implies that

$$\text{gap}(\rho) = \min_{(x,a) \in \rho} (f^*(x) - f_a(x)) \leq 4L\Delta(\rho_t) + 2\sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho_t)}}$$

Therefore if  $\rho$  is selected by the UCB rule at trial  $t$ , it must be that

$$n_t(\rho_t) \leq \frac{24\sigma^2 \ln(T)}{(\text{gap}(\rho) - 4L\Delta(\rho_t))^2}.$$

□

*Proof of Lemma B.5.* Note that  $\max_{a \in \mathcal{A}(\rho)} (f^*(x) - f_a(x))$  is a  $2L$ -Lipschitz function such that

$$\max_{x \in [c_0(\rho), c_1(\rho)]} \max_{a \in \mathcal{A}(\rho)} (f^*(x) - f_a(x)) \quad (30)$$

$$\leq \min_{x \in [c_0(\rho), c_1(\rho)]} \max_{a \in \mathcal{A}(\rho)} (f^*(x) - f_a(x)) + 2L\Delta(\rho) \quad (31)$$

$$\leq \min_{x \in [c_0(\rho), c_1(\rho)]} \min_{a \in \mathcal{A}(\rho)} (f^*(x) - f_a(x)) + \text{diam}(\rho) + 2L\Delta(\rho) \quad (32)$$

$$= \text{gap}(\rho) + \text{diam}(\rho) + 2L\Delta(\rho). \quad (33)$$

Conditioned on the “good events” holding, by Lemma B.4, if a ball  $\rho$  is flagged at trial  $t = \tau_f(\rho)$ , then it must be that

$$n_t(\rho) = \left\lfloor \frac{6\sigma^2 \ln(T)}{L^2 \Delta^2(\rho)} + 1 \right\rfloor \leq \frac{24\sigma^2 \ln(T)}{(\text{gap}(\rho) - 4L\Delta(\rho))^2}. \quad (34)$$

If the above inequality does not hold, then it would imply that the ball is so suboptimal that it stop being chosen by the UCB rule before hitting the threshold for flagging.

This implies

$$\frac{6\sigma^2 \ln(T)}{L^2 \Delta^2(\rho)} \leq \frac{24\sigma^2 \ln(T)}{(\text{gap}(\rho) - 4L\Delta(\rho))^2} \quad (35)$$

and thus

$$\text{gap}(\rho) \leq 6L\Delta(\rho). \quad (36)$$

By Lemma B.3 and (33), it follows that

$$\begin{aligned} \max_{(x,a) \in \rho} (f^*(x) - f_a(x)) &\leq \text{gap}(\rho) + \text{diam}(\rho) + 2L\Delta(\rho) \\ &\leq 6L\Delta(\rho) + 2L\Delta(\rho) + 2L\Delta(\rho) \\ &= 10L\Delta(\rho). \end{aligned}$$

This implies that the maximum regret for a ball who is eventually flagged is upper bounded by  $10L\Delta(\rho)$ . For any subsequent children balls  $\rho'$  that are formed by subpartitioning  $\rho$ , it must follow that

$$\text{gap}(\rho') \leq \max_{(x,a) \in \rho} (f^*(x) - f_a(x)) \quad (37)$$

$$\leq 10L\Delta(\rho) \quad (38)$$

$$= 20L\Delta(\rho') \quad (39)$$

□

*Proof of Lemma B.6.* The length of trial that a ball  $\rho$  stays in the flagged phase before triggering the SUBPARTITION subroutine depends on the number of samples collected for this ball until the condition  $\prod_{a \in \mathcal{A}(\rho)} \text{SUFFDATA}(a) == 1$ . Essentially this condition considers the 64 equal sized intervals that split the context space  $[c_0(\rho), c_1(\rho)]$  and checks that in each there is at least  $k$  points sampled for each arm  $a \in \mathcal{A}(\rho)$  within each of those 64 subintervals.

Due to the fact that the algorithm always gives priority to flagged balls, and furthermore flagged balls with larger context width are always given priority, it follows that the context of samples collected for ball  $\rho$  within its flagged phase must be distributed uniformly within  $[c_0(\rho), c_1(\rho)]$ . If  $\rho$  is the only ball flagged that intersects with the context interval  $[c_0(\rho), c_1(\rho)]$ , then it is given first priority such that any context that falls within this interval will be assigned to  $\rho$ . As the contexts arrive uniformly sampled over  $[0, 1]$ , the set of contexts restricted to  $[c_0(\rho), c_1(\rho)]$  will also be uniformly distributed.

As the context widths are always split into half, the endpoints must be equal to  $\ell 2^{-i}$  for some integers  $i$  and  $\ell$ . In the case that there is some other ball  $\rho'$  which intersects with  $[c_0(\rho), c_1(\rho)]$ , either it has smaller context width and thus must be fully contained within  $[c_0(\rho), c_1(\rho)]$ , or it has larger context width and must be a strict superset fully eclipsing  $[c_0(\rho), c_1(\rho)]$ . It is impossible for there to be another flagged ball  $\rho'$  with the exact same context width  $[c_0(\rho), c_1(\rho)]$ , because whichever ball was flagged first, would cause the algorithm to give full priority to that flagged ball, so that it would be impossible while that ball is still flagged, for any context in  $[c_0(\rho), c_1(\rho)]$  to be assigned to the other ball to trigger it to be flagged. If  $\rho'$  has smaller width, then  $\rho'$  is given lower priority, and no samples will be assigned to  $\rho'$  until  $\rho$  is subpartitioned and unflagged. If  $\rho'$  has larger width, then  $\rho'$  will be given higher priority over the entire interval  $[c_0(\rho), c_1(\rho)]$ , such that no samples will be assigned to  $\rho$  until  $\rho'$  is subpartitioned and unflagged.

As a result, in each trial  $t$ , either  $\rho$  has priority on the entire context interval  $[c_0(\rho), c_1(\rho)]$  and thus receives samples uniformly within that interval, or  $\rho$  does not have priority on any subset of the interval  $[c_0(\rho), c_1(\rho)]$  such that it receives no samples at all, guarantees that the eventual set of samples collected during the flagged phase must be distributed uniformly on  $[c_0(\rho), c_1(\rho)]$ . As a result, the probability that each sample collected for  $\rho$  falls into any of the 64 subintervals of  $[c_0(\rho), c_1(\rho)]$  is evenly  $1/64$ . By coupon collector, the expected number of samples until we get one sample in each bucket is  $64 \sum_{i=1}^{64} \frac{1}{i}$ . A naive upper bound on the number of samples until we get  $k$  samples in each bucket is  $64k \sum_{i=1}^{64} \frac{1}{i} \leq 304k$ .

As there are  $|\mathcal{A}(\rho)|$  arms, each of which needs to satisfy  $\text{SUFFDATA}(a) == 1$ ,

$$\mathbb{E} \left[ \sum_{t=\tau_f(\rho)+1}^T \mathbb{I}(\rho_t = \rho) \mid \tau_f(\rho) \leq T \right] \leq 304k|\mathcal{A}(\rho)| = \frac{304 \cdot 5431 \cdot \sigma^2 |\mathcal{A}(\rho)| \ln(T|\mathcal{A}(\rho)|)}{L^2 \Delta^2(\rho)} \quad (40)$$

□

## D Bounding Probability of Bad Events

*Proof of Lemma B.1.* Recall that by definition,

$$\begin{aligned}\mu_t(\rho) &= \frac{1}{n_t(\rho)} \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho) \pi_s \\ &= \frac{1}{n_t(\rho)} \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho) (f_{a_s}(x_s) + \epsilon_s)\end{aligned}$$

As  $(x_s, a_s) \in \rho$  if  $\rho_s = \rho$ , it follows that

$$\frac{1}{n_t(\rho)} \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho) f_{a_s}(x_s) \in \left[ \min_{(x,a) \in \rho} f_a(x), \max_{(x,a) \in \rho} f_a(x) \right].$$

It remains to be shown that with high probability, for all  $t$ ,

$$\left| \frac{1}{n_t(\rho)} \sum_{s=1}^{t-1} \mathbb{I}(\rho_s = \rho) \epsilon_s \right| \leq \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}}.$$

Note that we really only need to concern ourselves with values of  $t$  after which the ball has been chosen at least once. Otherwise, before the ball has been chosen yet, there are no terms to sum over, thus trivially zero is bounded above by the confidence bound.

For some ball  $\rho$ , let us denote the sequence

$$(\tau_1, \tau_2, \tau_3, \dots)$$

where  $\tau_s$  corresponds to the  $s$ -th trial that ball  $\rho$  is chosen by the algorithm, i.e.

$$\sum_{t=1}^{\tau_s} \mathbb{I}(\rho_t = \rho) = s \text{ and } \rho_{\tau_s} = \rho.$$

By Doob's optional skipping theorem,

$$(\epsilon_{\tau_1}, \epsilon_{\tau_2}, \epsilon_{\tau_3}, \dots)$$

is identically distributed to

$$(\epsilon_1, \epsilon_2, \epsilon_3, \dots).$$

Therefore, by Doob's optional skipping theorem, union bound, and Hoeffding's inequality,

$$\begin{aligned}& \mathbb{P}(\forall s \in [T], \left| \frac{1}{s} \sum_{\ell=1}^s \epsilon_{\tau_\ell} \right| \leq \sqrt{\frac{6\sigma^2 \ln(T)}{s}}) \\ &= \mathbb{P}(\forall s \in [T], \left| \frac{1}{s} \sum_{\ell=1}^s \epsilon_\ell \right| \leq \sqrt{\frac{6\sigma^2 \ln(T)}{s}}) \\ &\leq \sum_{s \in [T]} \mathbb{P}(\left| \sum_{\ell=1}^s \epsilon_\ell \right| \leq \sqrt{\frac{6s^2 \sigma^2 \ln(T)}{s}}) \\ &\leq \sum_{s \in [T]} 2 \exp(-3 \ln(T)) \\ &\leq 2T^{-2}\end{aligned}$$

There are at most  $T$  active balls over the course of the algorithm, thus by union bound over all active balls  $\rho$  over the course of the algorithm, with probability at least  $1 - 2T^{-1}$ , for all active balls  $\rho$ , for all  $t : n_t(\rho) \geq 1$ ,

$$\mu_t(\rho) \in \left[ \min_{(x,a) \in \rho} f_a(x) - \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}}, \max_{(x,a) \in \rho} f_a(x) + \sqrt{\frac{6\sigma^2 \ln(T)}{n_t(\rho)}} \right].$$

□

*Proof of Lemma B.2.* Lemma D.1 proves that each trial  $\hat{\mathcal{D}}_u^v(a, y)$  is evaluated within the subroutine SUBPARTITION( $\rho$ ) over the course of the algorithm,

$$\left\{ |\hat{\mathcal{D}}_u^v(a, y) - \mathcal{D}_u^v(a, y)| \leq \frac{1}{8}L(v-u) \right\}$$

with probability at least  $1 - 4T^{-2}|\mathcal{A}(\rho)|^{-2}$ .

Within a subroutine SUBPARTITION( $\rho$ ), the maximum number of trials that  $\hat{\mathcal{D}}_u^v(a, y)$  could be evaluated is  $|\mathcal{A}(\rho)|^{-2}$ . There are maximally  $T$  trials that subroutine SUBPARTITION can be called. By union bound, with probability  $1 - 4T^{-1}$ , over the entire course of the algorithm, for all trials that  $\hat{\mathcal{D}}_u^v(a, y)$  is evaluated,

$$\left\{ |\hat{\mathcal{D}}_u^v(a, y) - \mathcal{D}_u^v(a, y)| \leq \frac{1}{8}L(v-u) \right\}.$$

□

**Lemma D.1.** *Each trial  $\hat{\mathcal{D}}_u^v(a, y)$  is evaluated within the subroutine SUBPARTITION over the course of the algorithm,*

$$\left\{ |\hat{\mathcal{D}}_u^v(a, y) - \mathcal{D}_u^v(a, y)| \leq \frac{1}{8}L(v-u) \right\}$$

with probability at least  $1 - 4T^{-2}|\mathcal{A}(\rho)|^{-2}$ .

*Proof.* Consider a single call to the subroutine SUBPARTITION for a ball  $\rho$ . By construction, the subroutine is only called when  $\prod_{a \in \mathcal{A}(\rho)} \text{SUFFDATA}(a) == 1$ , where

$$\text{SUFFDATA}(a) = \prod_{i=1}^{64} \mathbb{I} \left( \sum_{s > \tau_f(\rho)} \mathbb{I}(\rho_s = \rho, a_s = a) \mathbb{I}(x_s \in [w_{i-1}, w_i]) \geq k \right) \quad (41)$$

for  $w_i = c_0(\rho) + \frac{i\Delta(\rho)}{64}$ . As a result, by construction, each trial  $\hat{\mathcal{D}}_u^v(a, y)$  is evaluated within the subroutine SUBPARTITION, if we take the interval  $[u, v]$  and split it evenly into 32 subintervals, the algorithm guarantees that for each of the 32 subintervals, for each arm  $a$  and  $y$ , there are at least  $k$  samples (or observations) collected for  $\rho$  during the “flagged phase” such that the context lies within the subinterval. As our algorithm estimates the reward functions  $\hat{f}_a$  and  $\hat{f}_y$  via  $k$  nearest neighbor averaging, this condition guarantees a minimum bias.

Recall that

$$\hat{\mathcal{D}}_u^v(a, y) := \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( \hat{f}_a(z_i(u, v)) - \hat{f}_y(z_i(u, v)) \right)^2 - \frac{2\sigma^2}{k}} \quad (42)$$

for

$$\hat{f}_a(x) = \frac{1}{k} \sum_{s=\tau_f(\rho)+1}^{\tau_{cl}(\rho)} \mathbb{I}(\rho_s = \rho, a_s = a) \mathbb{I}(x_s \in \text{k-NN}) \pi_s, \quad (43)$$

where  $x_s$  is a  $k$  nearest neighbor datapoint for computing  $\hat{f}_a(x)$  if

$$\mathbb{I}(x_s \in \text{k-NN}) = \sum_{\ell=\tau_f(\rho)+1}^{\tau_{cl}(\rho)} \mathbb{I}(\rho_\ell = \rho, a_\ell = a) \mathbb{I}(|x_\ell - x| \leq |x_s - x|) \leq k.$$

Recall that  $\pi_s = f_{a_s}(x_s) + \epsilon_s$ , where  $\epsilon_s$  is an independent noise term distributed as  $N(0, \sigma^2)$ . Let us denote

$$\bar{\mathcal{D}}_u^v(a, y) := \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( \bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v)) \right)^2} \quad (44)$$

for

$$\bar{f}_a(x) = \frac{1}{k} \sum_{s=\tau_f(\rho)+1}^{\tau_{cl}(\rho)} \mathbb{I}(\rho_s = \rho, a_s = a) \mathbb{I}(x_s \in \text{k-NN}) f_a(x_s). \quad (45)$$

As a result of the condition  $\prod_{a \in \mathcal{A}(\rho)} \text{SUFFDATA}(a) == 1$ , for any  $x$ , if  $x_s \in \text{k-NN}$ , then  $x_s - x \leq \frac{v-u}{32}$ , such that  $|f_a(x) - f_a(x_s)| \leq \frac{L(v-u)}{32}$ . This implies that  $|f_a(x) - \bar{f}_a(x)| \leq \frac{L(v-u)}{32}$ .

By triangle inequality and Lipschitzness,

$$\begin{aligned}
& |\mathcal{D}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| \\
& \leq \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( (f_a(z_i(u, v)) - f_y(z_i(u, v))) - (\bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v))) \right)^2} \\
& \leq \sqrt{\frac{1}{200} \sum_{i \in [200]} \left( \frac{L(v-u)}{16} \right)^2} \\
& \leq \frac{L(v-u)}{16}.
\end{aligned}$$

By Lemma D.2, with probability at least  $1 - 2T^{-2}|\mathcal{A}(\rho)|^{-2}$ ,  $|\bar{\mathcal{D}}_u^v(a, y) - \hat{\mathcal{D}}_u^v(a, y)| \leq \frac{L(v-u)}{16}$ . It follows that

$$|\mathcal{D}_u^v(a, y) - \hat{\mathcal{D}}_u^v(a, y)| \leq |\mathcal{D}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| + |\bar{\mathcal{D}}_u^v(a, y) - \hat{\mathcal{D}}_u^v(a, y)| \leq \frac{L(v-u)}{8}.$$

□

**Lemma D.2.** With probability at least  $1 - 4T^{-2}|\mathcal{A}(\rho)|^{-2}$ ,

$$|\hat{\mathcal{D}}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| \leq \frac{L(v-u)}{16}.$$

*Proof.* Let us denote  $\mathcal{E}(\rho, a, i) = \{t \in [\tau_f(\rho) + 1, \tau_{cl}(\rho)] \text{ s.t. } \rho_t = \rho, a_t = a, x_t \in \text{k-NN of } z_i(u, v)\}$ . Recall that by definition,

$$\hat{f}_a(z_i(u, v)) = \bar{f}_a(z_i(u, v)) + \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, a, i)} \epsilon_t$$

We split  $|(\hat{\mathcal{D}}_u^v(a, y))^2 - (\bar{\mathcal{D}}_u^v(a, y))^2|$  into two terms,

$$|(\hat{\mathcal{D}}_u^v(a, y))^2 - (\bar{\mathcal{D}}_u^v(a, y))^2| \tag{46}$$

$$= \left| \frac{1}{200} \sum_{i \in [200]} \left( (\hat{f}_a(z_i(u, v)) - \hat{f}_y(z_i(u, v)))^2 - (\bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v)))^2 \right) - \frac{2\sigma^2}{k} \right| \tag{47}$$

$$\leq \left| \frac{2}{200} \sum_{i \in [200]} (\bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v))) \left( \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, a, i)} \epsilon_t - \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, y, i)} \epsilon_t \right) \right| \tag{48}$$

$$+ \left| \frac{1}{200} \sum_{i \in [200]} \left( \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, a, i)} \epsilon_t - \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, y, i)} \epsilon_t \right)^2 - \frac{2\sigma^2}{k} \right| \tag{49}$$

By Lemmas D.3 and D.4, with probability at least  $1 - 4T^{-2}|\mathcal{A}(\rho)|^{-2}$ ,

$$|(\hat{\mathcal{D}}_u^v(a, y))^2 - (\bar{\mathcal{D}}_u^v(a, y))^2| \leq \bar{\mathcal{D}}_u^v(a, y) \sqrt{\frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}} + \frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}.$$

Let us denote  $Q_1 = \sqrt{\frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}}$ .

We consider two bounds for  $|\bar{\mathcal{D}}_u^v(a, y) - \hat{\mathcal{D}}_u^v(a, y)|$ ,

$$\begin{aligned}
|\hat{\mathcal{D}}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| & \leq \hat{\mathcal{D}}_u^v(a, y) + \bar{\mathcal{D}}_u^v(a, y) \\
& \leq \sqrt{|(\hat{\mathcal{D}}_u^v(a, y))^2 - (\bar{\mathcal{D}}_u^v(a, y))^2| + (\bar{\mathcal{D}}_u^v(a, y))^2} + \bar{\mathcal{D}}_u^v(a, y) \\
& \leq \sqrt{Q_1(\bar{\mathcal{D}}_u^v(a, y) + Q_1) + (\bar{\mathcal{D}}_u^v(a, y))^2} + \bar{\mathcal{D}}_u^v(a, y)
\end{aligned}$$

and

$$\begin{aligned}
|\hat{\mathcal{D}}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| & \leq \frac{|(\hat{\mathcal{D}}_u^v(a, y))^2 - (\bar{\mathcal{D}}_u^v(a, y))^2|}{\hat{\mathcal{D}}_u^v(a, y) + \bar{\mathcal{D}}_u^v(a, y)} \\
& \leq \frac{|(\hat{\mathcal{D}}_u^v(a, y))^2 - (\bar{\mathcal{D}}_u^v(a, y))^2|}{\bar{\mathcal{D}}_u^v(a, y)} \\
& \leq \frac{Q_1(\bar{\mathcal{D}}_u^v(a, y) + Q_1)}{\bar{\mathcal{D}}_u^v(a, y)}.
\end{aligned}$$

The first bound is strictly increasing in  $\bar{\mathcal{D}}_u^v(a, y)$ , whereas the second bound is strictly decreasing in  $\bar{\mathcal{D}}_u^v(a, y)$ . Therefore for any  $Q_2$ ,

$$\begin{aligned} |\hat{\mathcal{D}}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| &\leq \min \left( \bar{\mathcal{D}}_u^v(a, y) + \sqrt{Q_1(\bar{\mathcal{D}}_u^v(a, y) + Q_1) + (\bar{\mathcal{D}}_u^v(a, y))^2}, Q_1 + \frac{Q_1^2}{\bar{\mathcal{D}}_u^v(a, y)} \right) \\ &\leq \max \left( Q_2 + \sqrt{Q_1 Q_2 + Q_1^2 + Q_2^2}, Q_1 + \frac{Q_1^2}{Q_2} \right) \end{aligned}$$

Let us choose  $Q_2 = \frac{1+\sqrt{13}}{6}Q_1$  so that

$$\begin{aligned} |\hat{\mathcal{D}}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| &\leq \frac{7 + \sqrt{13}}{1 + \sqrt{13}} Q_1 \\ &\leq \frac{7 + \sqrt{13}}{1 + \sqrt{13}} \sqrt{\frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}} \end{aligned}$$

For

$$k = \frac{5431\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{L^2(v-u)^2} \geq 16^2 \left( \frac{7 + \sqrt{13}}{1 + \sqrt{13}} \right)^2 \frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{L^2(v-u)^2},$$

it follows that

$$|\hat{\mathcal{D}}_u^v(a, y) - \bar{\mathcal{D}}_u^v(a, y)| \leq \frac{L(v-u)}{16}.$$

□

**Lemma D.3.** *With probability at least  $1 - 2T^{-2}|\mathcal{A}(\rho)|^{-2}$ ,*

$$(48) \leq \bar{\mathcal{D}}_u^v(a, y) \sqrt{\frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}}.$$

*Proof.* Note that  $\frac{2}{200} \sum_{i \in [200]} (\bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v))) \left( \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, a, i)} \epsilon_t - \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, y, i)} \epsilon_t \right)$  is simply a mean zero Gaussian random variable, thus we only need to compute the variance and then we can apply Hoeffding's Inequality.

Let us denote

$$Y_i = (\bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v))) \left( \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, a, i)} \epsilon_t - \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, y, i)} \epsilon_t \right),$$

such that the expression of interest to us can be restated as  $\frac{2}{200} \sum_{i \in [200]} Y_i$ . Note that

$$\text{Var}[Y_i] = (\bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v)))^2 \frac{2\sigma^2}{k}$$

Then

$$\text{Var}\left[\frac{2}{200} \sum_{i \in [200]} Y_i\right] = \frac{4}{200^2} \sum_{i, i' \in [200]^2} \text{Cov}(Y_i, Y_{i'})$$

By construction,  $\text{Cov}(Y_i, Y_{i'})$  is zero if  $|z_i(u, v) - z_{i'}(u, v)| > \frac{v-u}{16}$ , as all the  $k - NN$  points must be within distance  $\frac{v-u}{32}$ . Additionally,

$$\text{Cov}(Y_i, Y_{i'}) \leq \frac{1}{2} (\text{Var}[Y_i] + \text{Var}[Y_{i'}]).$$

Then

$$\begin{aligned}
& \frac{4}{200^2} \sum_{i,i' \in [200]^2} \text{Cov}(Y_i, Y_{i'}) \\
& \leq \frac{4}{200^2} \sum_{i,i' \in [200]^2} \frac{1}{2} (\text{Var}[Y_i] + \text{Var}[Y_{i'}]) \mathbb{I}(|z_i(u, v) - z_{i'}(u, v)| \leq \frac{v-u}{16}) \\
& \leq \frac{4}{200^2} \sum_{i \in [200]} \text{Var}[Y_i] \sum_{i' \in [200]} \mathbb{I}(|z_i(u, v) - z_{i'}(u, v)| \leq \frac{v-u}{16}) \\
& \leq \frac{4}{200^2} \sum_{i \in [200]} \text{Var}[Y_i] \max(1, \frac{200}{8}) \\
& \leq \frac{\sigma^2}{200k} \sum_{i \in [200]} (\bar{f}_a(z_i(u, v)) - \bar{f}_y(z_i(u, v)))^2 \\
& = \frac{\sigma^2}{k} (\bar{\mathcal{D}}_u^v(a, y))^2
\end{aligned}$$

From Hoeffding's inequality,

$$P[(48) \geq \delta] \leq 2 \exp \left\{ - \frac{\delta^2 k}{2\sigma^2 (\bar{\mathcal{D}}_u^v(a, y))^2} \right\} \quad (50)$$

For  $\delta = \bar{\mathcal{D}}_u^v(a, y) \sqrt{\frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}}$ , with probability at least  $1 - 2T^{-2}|\mathcal{A}(\rho)|^{-2}$ ,  $(48) < \delta$ .  $\square$

**Lemma D.4.** *With probability at least  $1 - 2T^{-2}|\mathcal{A}(\rho)|^{-2}$ ,*

$$(49) \leq \frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}.$$

*Proof.* Let us define

$$Y_i = \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, a, i)} \epsilon_t - \frac{1}{k} \sum_{t \in \mathcal{E}(\rho, y, i)} \epsilon_t,$$

so that the expression in (49) can be rewritten as  $|\frac{1}{200} \sum_{i \in [200]} Y_i^2 - \frac{2\sigma^2}{k}|$ . The left and right terms of  $Y_i$  are independent as they correspond to samples obtained for different arms,  $a$  and  $y$ , so  $\mathcal{E}(\rho, a, i)$  is completely disjoint from  $\mathcal{E}(\rho, y, i)$ . Therefore  $Y_i \sim N(0, \frac{2\sigma^2}{k})$  such that  $\mathbb{E} \left[ \frac{1}{200} \sum_{i \in [200]} Y_i^2 \right] = \frac{2\sigma^2}{k}$ .

Next, we want to show that  $\frac{1}{200} \sum_{i \in [200]} Y_i^2$  is sub-exponential, so that we can use Bernstein's inequality to bound its concentration around its mean. The vector  $(Y_i)_{i \in [200]}$  can be written as a affine transformation  $QX$ , where  $Q$  is the matrix defined as

$$Q_{it} = \frac{1}{k} (\mathbb{I}(t \in \mathcal{E}(\rho, a, i)) - \mathbb{I}(t \in \mathcal{E}(\rho, y, i))),$$

and  $X_t = \epsilon_t$ , such that  $X \sim N(0, \sigma^2 I)$ . Therefore the vector  $(Y_i)_{i \in [200]}$  is a multivariate Gaussian with mean zero and variance  $\sigma^2 Q Q^T$ .

It holds by the sub-exponential property of sum of Gaussians that conditioned on the latent variables  $\{\beta_i\}_{i \in [m]}$  and the observation indices  $\mathcal{E}'$ ,

$$\mathbb{P} \left( \left| \frac{Y^T Y}{200} - \frac{\|\sigma^2 Q Q^T\|_*}{200} \right| \geq \delta \right) \leq \begin{cases} 2 \exp \left( - \frac{200^2 \delta^2}{8 \|\sigma^2 Q Q^T\|_F^2} \right) & \text{if } \delta \leq \frac{\|\sigma^2 Q Q^T\|_F^2}{200 \|\sigma^2 Q Q^T\|_2} \\ 2 \exp \left( - \frac{200 \delta}{8 \|\sigma^2 Q Q^T\|_2} \right) & \text{if } \delta > \frac{\|\sigma^2 Q Q^T\|_F^2}{200 \|\sigma^2 Q Q^T\|_2} \end{cases}.$$

$\|\cdot\|_*$  denotes the nuclear norm,  $\|\cdot\|_F$  denotes the frobenius norm, and  $\|\cdot\|_2$  denotes the spectral norm. We can verify that indeed

$$\frac{\|\sigma^2 Q Q^T\|_*}{200} = \frac{\sigma^2 \text{Tr}(Q Q^T)}{200} = \frac{2\sigma^2}{k}.$$

Let us first upper bound  $[[QQ^T]_{ij}]$ ,

$$\begin{aligned}
|[QQ^T]_{ij}| &= \left| \sum_t Q_{it} Q_{jt} \right| \\
&= \left| \frac{1}{k^2} \sum_t (\mathbb{I}(t \in \mathcal{E}(\rho, a, i)) - \mathbb{I}(t \in \mathcal{E}(\rho, y, i))) (\mathbb{I}(t \in \mathcal{E}(\rho, a, j)) - \mathbb{I}(t \in \mathcal{E}(\rho, y, j))) \right| \\
&\leq \mathbb{I} \left( |z_i(u, v) - z_j(u, v)| \leq \frac{v-u}{16} \right) \frac{2}{k}.
\end{aligned}$$

We use this to upper bound  $\|\sigma^2 QQ^T\|_F^2$ ,

$$\begin{aligned}
\|\sigma^2 QQ^T\|_F^2 &= \sum_{i \in [200]} \sum_{j \in [200]} ([\sigma^2 QQ^T]_{ij})^2 \\
&\leq \sigma^4 \sum_{i \in [200]} \sum_{j \in [200]} \mathbb{I} \left( |z_i(u, v) - z_j(u, v)| \leq \frac{v-u}{16} \right) \frac{4}{k^2} \\
&\leq \frac{4\sigma^4}{k^2} \sum_{i \in [200]} \frac{200}{8} \\
&= \frac{200^2 \sigma^4}{2k^2}
\end{aligned}$$

By symmetry,  $\|\sigma^2 QQ^T\|_1 = \|\sigma^2 QQ^T\|_\infty$ . By Holder's inequality,

$$\|\sigma^2 QQ^T\|_2 \leq \sqrt{\|\sigma^2 QQ^T\|_1 \|\sigma^2 QQ^T\|_\infty} \quad (51)$$

$$= \max_{j \in [200]} \sum_{i \in [200]} |[\sigma^2 QQ^T]_{ij}| \quad (52)$$

$$\leq \frac{2\sigma^2}{k} \max_{j \in [200]} \sum_{i \in [200]} \mathbb{I} \left( |z_i(u, v) - z_j(u, v)| \leq \frac{v-u}{16} \right) \quad (53)$$

$$\leq \frac{2\sigma^2}{k} \max_{j \in [200]} \frac{200}{8} \quad (54)$$

$$\leq \frac{200\sigma^2}{4k} \quad (55)$$

By plugging these bounds in, we obtain that

$$\mathbb{P} \left( \left| \frac{Y^T Y}{200} - \frac{2\sigma^2}{k} \right| \geq \delta \right) \leq \begin{cases} 2 \exp \left( -\frac{\delta^2 k^2}{4\sigma^4} \right) & \text{if } \delta \leq \frac{2\sigma^2}{k} \\ 2 \exp \left( -\frac{\delta k}{2\sigma^2} \right) & \text{if } \delta > \frac{2\sigma^2}{k} \end{cases}.$$

Let us choose  $\delta = \frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}$ , such that  $\delta > \frac{2\sigma^2}{k}$ . Therefore, with probability greater than  $1 - 2T^{-2}|\mathcal{A}(\rho)|^{-2}$ , (49)  $\leq \frac{4\sigma^2 \ln(T|\mathcal{A}(\rho)|)}{k}$ .

□

## E Final Regret Calculation

*Proof of Theorem 6.1.* When we sum over  $\rho \in \mathcal{P}$ , we mean to refer to all balls over trial that are ever active, i.e. a member of  $\mathcal{P}$  at any point of the algorithm within the time horizon  $T$ .

By using a similar argument to Lemma B.6, the regret from initial clustering is bounded above by

$$\frac{304 \cdot 5431 \cdot \sigma^2 K \ln(TK)}{L^2} = O \left( \frac{\sigma^2 K \ln(TK)}{L^2} \right),$$



where we used the property that regret is bounded above by 1 in every trial step because the expected reward function  $f$  outputs values in  $[0, 1]$ .

Next we bound the expected regret after the initial clustering conditioned on the good events  $\mathcal{G}$ . We split the regret into two terms,

$$\mathbb{E}[\text{regret after initial clustering} \mid \mathcal{G}] \quad (56)$$

$$= \mathbb{E} \left[ \sum_{t=1}^T \sum_{\rho \in \mathcal{P}} \mathbb{I}(\rho_t = \rho) (f^*(x_t) - f_{a_t}(x_t)) \mid \mathcal{G} \right] \quad (57)$$

$$= \mathbb{E} \left[ \sum_{t=1}^T \sum_{i=1}^{i_{\max}-1} \sum_{\rho \in \mathcal{P}} \mathbb{I}(\rho_t = \rho, \Delta(\rho) = 2^{-i}) (f^*(x_t) - f_{a_t}(x_t)) \mid \mathcal{G} \right] \quad (58)$$

$$+ \mathbb{E} \left[ \sum_{t=1}^T \sum_{i=i_{\max}}^{\infty} \sum_{\rho \in \mathcal{P}} \mathbb{I}(\rho_t = \rho, \Delta(\rho) = 2^{-i}) (f^*(x_t) - f_{a_t}(x_t)) \mid \mathcal{G} \right] \quad (59)$$

By Lemma B.5, (59) is bounded above by

$$\mathbb{E} \left[ 20L2^{-i_{\max}} \sum_{t=1}^T \sum_{i=i_{\max}}^{\infty} \sum_{\rho \in \mathcal{P}} \mathbb{I}(\rho_t = \rho, \Delta(\rho) = 2^{-i}) \mid \mathcal{G} \right] \leq 20L2^{-i_{\max}}T$$

Due to the fact that our algorithm always halved the context width, any ball  $\rho$  with  $\Delta(\rho) = 2^{-i}$  must have context width  $w_i(\ell) = [(\ell-1)2^{-i}, \ell 2^{-i}]$  for some  $\ell \in [2^i]$ . As a result, (58) equals

$$\mathbb{E} \left[ \sum_{i=1}^{i_{\max}-1} \sum_{\ell=1}^{2^i} \sum_{\rho \in \mathcal{P}} \mathbb{I}([c_0(\rho), c_1(\rho)] = w_i(\ell)) \sum_{t=1}^T \mathbb{I}(\rho_t = \rho) (f^*(x_t) - f_{a_t}(x_t)) \mid \mathcal{G} \right]$$

By Lemma B.5, conditioned on the good events  $\mathcal{G}$ , for any  $\rho \in \mathcal{P}$  such that  $\Delta(\rho) \leq \frac{1}{4}$ ,

$$\max_{(x,a) \in \rho} (f^*(x) - f_a(x)) \leq 20L\Delta(\rho).$$

This implies that for a context width  $c_0(\rho), c_1(\rho)$ , we can eliminate all arms  $a \in [K]$  that are “extremely suboptimal”, satisfying

$$\min_{x \in [c_0(\rho), c_1(\rho)]} (f^*(x) - f_a(x)) > 20L\Delta(\rho).$$

Let  $\kappa(x)$  denote the suboptimality gap at context  $x$ , i.e. the difference between the optimal set of arms and the next optimal set of arms,

$$\kappa(x) = f^*(x) - \sup_{a \in [K]} f_a(x) \mathbb{I}(f_a(x) \neq f^*(x)).$$

If  $\kappa(x) > 20L2^{-i}$  for all  $x \in w_i(\ell)$ , Lemma B.5 implies that conditioned on the good event  $\mathcal{G}$ , the regret incurred by any ball  $\rho$  for which  $[c_0(\rho), c_1(\rho)] \subseteq w_i(\ell)$  must be zero, as it must contain only optimal arms.

Thus we can reduce (58) to

$$\sum_{i=1}^{i_{\max}-1} \sum_{\ell=1}^{2^i} \mathbb{I} \left( \min_{x \in w_i(\ell)} \kappa(x) \leq 20L2^{-i} \right) \mathbb{E} \left[ \sum_{\rho \in \mathcal{P}} \mathbb{I}([c_0(\rho), c_1(\rho)] = w_i(\ell)) \sum_{t=1}^T \mathbb{I}(\rho_t = \rho) (f^*(x_t) - f_{a_t}(x_t)) \mid \mathcal{G} \right]$$

By Lemmas B.5, B.4, and B.6, and using a trivial upper bound that  $|\mathcal{A}(\rho)| \leq K$ ,

$$\begin{aligned}
& \mathbb{E} \left[ \sum_{\rho \in \mathcal{P}} \mathbb{I}([c_0(\rho), c_1(\rho)] = w_i(\ell)) \left( \sum_{t=1}^{\tau_f(\rho)} \mathbb{I}(\rho_t = \rho) (f^*(x_t) - f_{a_t}(x_t)) + \sum_{t=\tau_f(\rho)+1}^T \mathbb{I}(\rho_t = \rho) (f^*(x_t) - f_{a_t}(x_t)) \right) \middle| \mathcal{G} \right] \\
& \leq \mathbb{E} \left[ \sum_{\rho \in \mathcal{P}} \mathbb{I}([c_0(\rho), c_1(\rho)] = w_i(\ell)) \left( 20L\Delta(\rho) \left( \frac{6\sigma^2 \ln(T)}{L^2 \Delta^2(\rho)} + 2 \right) + 10L\Delta(\rho) \left( \frac{304 \cdot 5431 \cdot \sigma^2 |\mathcal{A}(\rho)| \ln(T|\mathcal{A}(\rho)|)}{L^2 \Delta^2(\rho)} \right) \right) \middle| \mathcal{G} \right] \\
& = O \left( \frac{\sigma^2 \ln(TK)}{L2^{-i}} \mathbb{E} \left[ \sum_{\rho \in \mathcal{P}} \mathbb{I}([c_0(\rho), c_1(\rho)] = w_i(\ell)) |\mathcal{A}(\rho)| \middle| \mathcal{G} \right] \right)
\end{aligned}$$

The algorithm guarantees that the active balls always form a partition of the context-arm space, and the balls are strictly nested in a hierarchy; as a result, the arms in different balls  $\rho \in \mathcal{P}$  constrained to the same context width must be disjoint.

By Lipschitzness,

$$\min_{x \in w_i(\ell)} (f^*(x) - f_a(x)) \leq 20L2^{-i} \implies \max_{x \in w_i(\ell)} (f^*(x) - f_a(x)) \leq 22L2^{-i},$$

such that

$$\mathbb{I} \left( \min_{x \in w_i(\ell)} (f^*(x) - f_a(x)) \leq 20L2^{-i} \right) \leq \mathbb{I}((f^*(2^{-i}\ell) - f_a(2^{-i}\ell)) \leq 22L2^{-i}).$$

For  $i \geq 2$ , by Lemma B.5, conditioned on the good events  $\mathcal{G}$ ,

$$\begin{aligned}
\mathbb{E} \left[ \sum_{\rho \in \mathcal{P}} \mathbb{I}([c_0(\rho), c_1(\rho)] = w_i(\ell)) |\mathcal{A}(\rho)| \middle| \mathcal{G} \right] & \leq \sum_{a \in [K]} \mathbb{I} \left( \min_{x \in w_i(\ell)} (f^*(x) - f_a(x)) \leq 20L2^{-i} \right) \\
& \leq \sum_{a \in [K]} \mathbb{I}((f^*(2^{-i}\ell) - f_a(2^{-i}\ell)) \leq 22L2^{-i}).
\end{aligned}$$

which upper bounds the number of arms that are ever subpartitioned into a ball of context width  $w_i(\ell)$ .

Let us denote

$$M_i = \sum_{\ell=1}^{2^i} \mathbb{I} \left( \min_{x \in w_i(\ell)} \kappa(x) \leq 20L2^{-i} \right) \sum_{a \in [K]} \mathbb{I}((f^*(2^{-i}\ell) - f_a(2^{-i}\ell)) \leq 22L2^{-i}).$$

The scaling of this quantity depends on the local geometry amongst the arms with respect to the expected reward functions.

To get the final regret bound, we use Lemmas B.1 and B.1 to bound the probability that the good event  $\mathcal{G}$  is violated,

$$\mathbb{E}[R(T)] \tag{60}$$

$$\leq \mathbb{E}[\text{regret from initial clustering}] + T\mathbb{P}(\neg\mathcal{G}) + \mathbb{E}[\text{regret after initial clustering} \mid \mathcal{G}] \tag{61}$$

$$\leq O \left( \frac{\sigma^2 K \ln(TK)}{L^2} \right) + T(2T^{-1} + 4T^{-1}) + 20LT2^{-i_{\max}} \tag{62}$$

$$+ O \left( \sum_{i=1}^{i_{\max}-1} \frac{\sigma^2 M_i \ln(TK)}{L2^{-i}} \right) \tag{63}$$

$$= O \left( \frac{\sigma^2 K \ln(TK)}{L^2} + \min_{i_{\max} \in \mathbb{Z}_+} \left( LT2^{-i_{\max}} + \sum_{i=1}^{i_{\max}-1} \frac{\sigma^2 M_i \ln(TK)}{L2^{-i}} \right) \right) \tag{64}$$

□

**Finite Types** Suppose that the reward functions for the  $K$  arms,  $\{f_a\}_{a \in [K]}$  only takes  $M$  different values. Essentially, this implies that there are  $\Theta$  different types of arms, but we don't know the arm types a priori. Within each type, the reward function is exactly the same. Let us denote the type of arm  $a$  with  $\theta_a \in [\Theta]$ , and we define function  $g : [0, 1] \times [\Theta] \rightarrow [0, 1]$  such that  $f_a(x) = g(x, \theta)$ . Let us define

$$\mu_\kappa(z) := \mu(\{x \in [0, 1] \text{ s.t. } \kappa(x) \leq z\})$$

where  $\mu$  is the Lebesgue measure.

Naively,  $\sum_{a \in [K]} \mathbb{I}((f^*(2^{-i}\ell) - f_a(2^{-i}\ell)) \leq 22L2^{-i}) \leq K$ . Furthermore, it will minimally be equal to the number of arms of the optimal type. If there are roughly proportional number of arms in each type, then this expression is also lower bounded by order  $K$ , thus it is sufficient to upper bound the expression by  $K$ . Then  $M_i$  can be upper bounded by

$$M_i = K \sum_{\ell=1}^{2^i} \mathbb{I}\left(\min_{x \in w_i(\ell)} \kappa(x) \leq 20L2^{-i}\right) \leq \frac{K\mu_\kappa(22L2^{-i})}{2^{-i}},$$

which follows from the fact that  $\kappa(x)$  must be a  $2L$ -Lipschitz function.

Let us denote the type of arm  $a$  with  $\theta_a \in [\Theta]$ , and we define function  $g : [0, 1] \times [\Theta] \rightarrow [0, 1]$  such that  $f_a(x) = g(x, \theta)$ . Let  $\theta^*(x)$  denote the set of arm types that are optimal at context  $x$ ,

$$\theta^*(x) = \{\theta \in [\Theta] \text{ s.t. } g(x, \theta) = f^*(x)\}$$

In the finite types setting, the optimal policy corresponds to partitioning the context space  $[0, 1]$  into a set of intervals,  $\mathcal{S}^*$ , such that across each interval  $f \in \mathcal{S}^*$ , the optimal policy does not change, i.e.  $\theta^*(x) = \theta^*(x')$  for all  $(x, x') \in f \times f$ . Note that at each endpoint of  $f$ , it must be that  $\kappa(x) = 0$ , as the fact that the policy changes and the reward functions are Lipschitz will imply that either an optimal arm becomes suboptimal, or a suboptimal arm becomes optimal, but the change must happen ‘‘smoothly’’ due to Lipschitzness. This also implies that for some  $\gamma$  arbitrarily close to 0, if  $x$  is an endpoint of any interval, either  $\kappa(x + \gamma) > 0$  or  $\kappa(x - \gamma) > 0$ .

Let us assume that  $\kappa(x)$  decreases linearly fast nearby the points where the optimal policy changes, so that for some constant  $L'$ ,

$$\mu_\kappa(22L2^{-i}) \leq \frac{22L2^{-i}}{L'} \times |\mathcal{S}^*|.$$

Then it follows by plugging into the main theorem that the regret is upper bounded by

$$\mathbb{E}[R(T)] = O\left(\frac{\sigma^2 K \ln(TK)}{L^2} + \min_{i_{\max} \in \mathbb{Z}_+} \left(LT2^{-i_{\max}} + \sum_{i=1}^{i_{\max}-1} \frac{22\sigma^2 |\mathcal{S}^*| K \ln(TK)}{L'2^{-i}}\right)\right). \quad (65)$$

By choosing

$$i_{\max} = \frac{1}{2} \log\left(\frac{L'LT}{22\sigma^2 |\mathcal{S}^*| K \ln(TK)}\right),$$

it follows that

$$\mathbb{E}[R(T)] \leq O\left(\frac{\sigma^2 K \ln(TK)}{L^2} + \sqrt{\frac{\sigma^2 |\mathcal{S}^*| LTK \ln(TK)}{L'}}\right) \quad (66)$$

**Lipschitz arm geometry** Suppose that each arm  $a$  is associated to a latent variable  $\theta_a \in [0, 1]$ , and the expected reward function  $f_a(x) = g(x, \theta_a)$ , where  $g : [0, 1] \times [0, 1] \rightarrow [0, 1]$  is a  $L$ -Lipschitz function with respect to both the contexts and the arm latent variables,

$$g(x, \theta) - g(x', \theta') \leq L(|x - x'| + |\theta - \theta'|).$$

If we assume that the arm latent variables are uniformly spread out,  $\{\theta_a\} = \{\frac{i}{K}\}_{i \in [K]}$ , then

$$M_i = \sum_{\ell=1}^{2^i} \mathbb{I}\left(\min_{x \in w_i(\ell)} \kappa(x) \leq 20L2^{-i}\right) \sum_{a \in [K]} \mathbb{I}((f^*(2^{-i}\ell) - g(2^{-i}\ell, \theta_a)) \leq 22L2^{-i}) \quad (67)$$

$$\leq \sum_{j \in [K]} \sum_{\ell=1}^{2^i} \mathbb{I}\left((f^*(2^{-i}\ell) - g(2^{-i}\ell, \frac{j}{K})) \leq 22L2^{-i}\right), \quad (68)$$

which is a discrete approximation to the area of the arm-context space for which the suboptimality gap is at most  $22L2^{-i}$ . We can visualize  $\sum_{\ell=1}^{2^i} M_i(\ell)$  by considering the contour plot of  $f^*(x) - g(x, \theta)$ , and counting how many grid points  $\{(2^{-i}\ell, \frac{j}{K})\}_{\ell \in [2^i], j \in [K]}$  are lower than  $22L2^{-i}$ . For large  $i$  and  $K$ , this is approximately equal to  $2^i K \mu(\{(x, \theta) : g(x, \theta) - f^*(x) \geq -22L2^{-i}\})$ , where  $\mu$  is the Lebesgue measure. The curve at the lowest level of the contour plot corresponds to the set  $\{(x, \theta) \text{ s.t. } g(x, \theta) - f^*(x) = 0\}$ , which contains for each context  $x$  the set of arm latent variables  $\theta$  that optimize the expected reward function. The final regret bound thus depends on the local measure/smoothness of the joint reward function.

To give a concrete example, we compute a bound for the reward function used in the simulation, where  $g(x, \theta) = 1 - L|x - \theta|$  for some  $L \in (0, 1)$ .

$$\begin{aligned}
M_i &\leq \sum_{\ell=1}^{2^i} \sum_{a \in [K]} \mathbb{I}((f^*(2^{-i}\ell) - f_a(2^{-i}\ell)) \leq 22L2^{-i}) \\
&= \sum_{\ell=1}^{2^i} \sum_{j \in [K]} \mathbb{I}\left((f^*(2^{-i}\ell) - g(2^{-i}\ell, \frac{j}{K})) \leq 22L2^{-i}\right) \\
&\leq \sum_{\ell=1}^{2^i} \sum_{j \in [K]} \mathbb{I}\left(L|2^{-i}\ell - \frac{j}{K}| \leq 22L2^{-i}\right) \\
&\leq \sum_{\ell=1}^{2^i} \sum_{j \in [K]} \mathbb{I}(j \in [K(2^{-i}\ell - 22 \cdot 2^{-i}), K(2^{-i}\ell + 22 \cdot 2^{-i})]) \\
&\leq \sum_{\ell=1}^{2^i} 44 \cdot 2^{-i} K \\
&\leq 44K
\end{aligned}$$

By plugging this into the main theorem, it follows that

$$\mathbb{E}[R(T)] = O\left(\frac{\sigma^2 K \ln(TK)}{L^2} + \min_{i_{\max} \in \mathbb{Z}_+} \left(LT2^{-i_{\max}} + \sum_{i=1}^{i_{\max}-1} \frac{\sigma^2 K \ln(TK)}{L2^{-i}}\right)\right). \quad (69)$$

Choosing

$$i_{\max} = \frac{1}{2} \log\left(\frac{20L^2 T}{\sigma^2 K \ln(TK)}\right),$$

results in

$$\mathbb{E}[R(T)] \leq O\left(\frac{\sigma^2 K \ln(TK)}{L^2} + \sqrt{\sigma^2 K T \ln(TK)}\right). \quad (70)$$

## F Additional Simulation Results and Discussion

We test our algorithm on a model with 50, 100, 200 arms and a context space of  $[0, 1]$ . Each arm  $a$  corresponds to a parameter  $\theta_a$  uniformly spaced out within  $[0, 1]$ . The expected reward for arm  $a$  and context  $x$  is

$$f_a(x) := g(x, \theta_a) = 1 - |x - 4 \min_{z \in \{0, 0.5, 1\}} |\theta_a - z||.$$

This function is periodic with respect to  $\theta$ , and can be depicted as a zigzag. Our distance estimate  $\hat{D}_u^v(a, a')$  approximates  $\mathcal{D}_u^v(a, a')$ , which is defined with respect  $f_a$  and  $f_{a'}$  directly and does not depend on  $\theta_a$ . Consider a measure preserving transformation that maps  $\theta_a$  to  $\phi_a = 4 \min_{z \in \{0, 0.5, 1\}} |\theta_a - z|$ , such that the reward function is equivalently described by  $f_a(x) = 1 - |x - \phi_a|$ . An algorithm which partitions with respect to  $\mathcal{D}_u^v(a, a')$  would be agnostic to such a transformation, as opposed to an algorithm which depends on a metric defined with respect to the arm's representation, which would perform worse on  $\theta_a$  than  $\phi_a$ .

In a sequence of  $T$  trials we uniformly randomly sample from the context space and reveal it to the algorithm. The algorithm selects what it considers to be the best arm in each trial based on the context revealed. Then simulation reveals a noisy payoff (i.e.,  $f_a(x) + \sigma$ ) to the algorithm. The task in the simulation setup<sup>1</sup> is for our algorithm to learn the optimal arm for different contexts.

We benchmark the performance of our Approx-Zooming algorithm against three variations:

- *Approx-Zooming -With-True-Reward-Function*: We give the Approx-Zooming algorithm oracle access to evaluate  $\mathcal{D}_u^v(a, a')$  at no cost, which is used to subpartition whenever a ball is flagged.
- *Approx-Zooming -With-Similarity-Metric*: We give the Approx-Zooming algorithm oracle access to evaluate  $|\theta_a - \theta_{a'}|$  at no cost, which is used to subpartition whenever a ball is flagged.
- *Approx-Zooming -With-No-Arm-Similarity*: This naive variant uses no arm similarities, estimating each arm’s reward independently. The context space is adaptively partitioned via our algorithm.

The possible combinations of the parameters we have chosen are as follows. For all algorithms we have evaluated for  $[50, 100, 200]$  arms over  $[10, 000, 50, 000, 100, 000]$  trials for  $\sigma$  values  $[1e - 1, 1e - 2, 1e - 3]$ . We chose the model parameters that led to the highest average cumulative reward in each baseline algorithm based on a selected set of parameter permutations. There were no hyper parameters to be considered for Approx-Zooming-No-Arm-Similarity. As per the implementation of the algorithms Approx-Zooming-Similarity-Metric and Approx-Zooming-True we had to specify a starting bias (from 0.25, 0.5, 1), starting distance threshold (from 0.25, 0.5, 1, 10) within a partition and the context discretization mapping (from 5, 6, 7). In addition to the above parameters as per the implementation of Approx-Zooming algorithm we had to specify the number of neighbours. The final  $k$  nearest neighbours for the Approx-Zooming algorithm was a function of a minimum value we specified (from 3, 4, 5), a global sampling constant (from 1, 2, 3), the number of neighbours we specified explicitly (from 5, 10, 15) and the context width. We evaluated the algorithms for some of these permutations and selected the best hyper-parameter configuration<sup>2</sup>. For all algorithms the flagging rule is set to  $n_t(\rho) \geq 4 \ln(T)/\Delta^2$ , and for results reported in this paper  $\sigma$  is set to either  $1e - 3$  or  $1e - 2$ . For Approx-Zooming,  $k$  was set to 10. We set the number of trials  $T$  to 100, 000 as all the algorithms had converged to their optimal point by then. We present results for the three simulation settings: (1) 50 arms with noise  $\sigma = 1e - 3$ , (2) 100 arms with  $\sigma = 1e - 3$  and (3) 200 arms with  $\sigma = 1e - 2$ .

In figure 3, we plot the average cumulative reward over the trials, i.e.  $\frac{1}{T} \sum_{t=1}^T \pi_t$ , where  $T$  is the total number of trials and  $\pi_t \in (0, 1)$  is the reward observed in the  $t^{\text{th}}$  trial for different simulation settings. As we can see, the oracle variant of the algorithm that uses the true reward function to calculate  $\mathcal{D}_u^v(a, a')$  performs the best on all three plots. Our Approx-Zooming algorithm has a heavy cost up front due to the clustering of the arms globally, but the algorithm improves over the time horizon as it learns the correct arm similarities. The oracle variant which uses the similarity metric  $|\theta_a - \theta_{a'}|$  performs worse than the true  $\mathcal{D}_u^v(a, a')$  variant, as it does not account for the periodic nature of the function.

In figure 4, figure 5 and figure 6 we plot the frequencies an arm is selected in different contexts over the  $T$  trials in our three simulation settings. Each of the four plots correspond to averaging the frequency over  $T/4$  trials across the time horizon. The x-axis refers to the context space, and the y-axis refers to the set of arms. As we see, the algorithms do generally learn to play the optimal policy, which corresponds to the zigzag shape. We can verify that in our algorithm initially the frequency plot is very blurry, indicating that it is spreading out the samples over time. As time progresses our algorithm indeed learns the similarities, which is depicted by the shape sharpening. The Approx-Zooming-True algorithm is given the true distance function, and we can see that indeed it is the sharpest curve. Approx-Zooming-Similarity-Metric, which is using the metric representation narrows in slower than the algorithm given true distance function. Approx-Zooming-No-Arm-Similarity which learns each arm separately initially finds a small set of arms that plays across the context and as a result takes more time to converge to the optimal policy.

The algorithm which learns each arm separately takes more time to converge to the optimal policy compared to all the other methods. Therefore, we can conclude that given a large arm set it is

<sup>1</sup>The code to the algorithms and simulation is available at [https://bitbucket.org/nirandiw/context\\_similarity/src/master/](https://bitbucket.org/nirandiw/context_similarity/src/master/)

<sup>2</sup>The hyper parameter tuning results are available at [https://drive.google.com/open?id=1vb-RK8E\\_f1PZ7haN\\_M83Y63Fes87j6fs](https://drive.google.com/open?id=1vb-RK8E_f1PZ7haN_M83Y63Fes87j6fs)

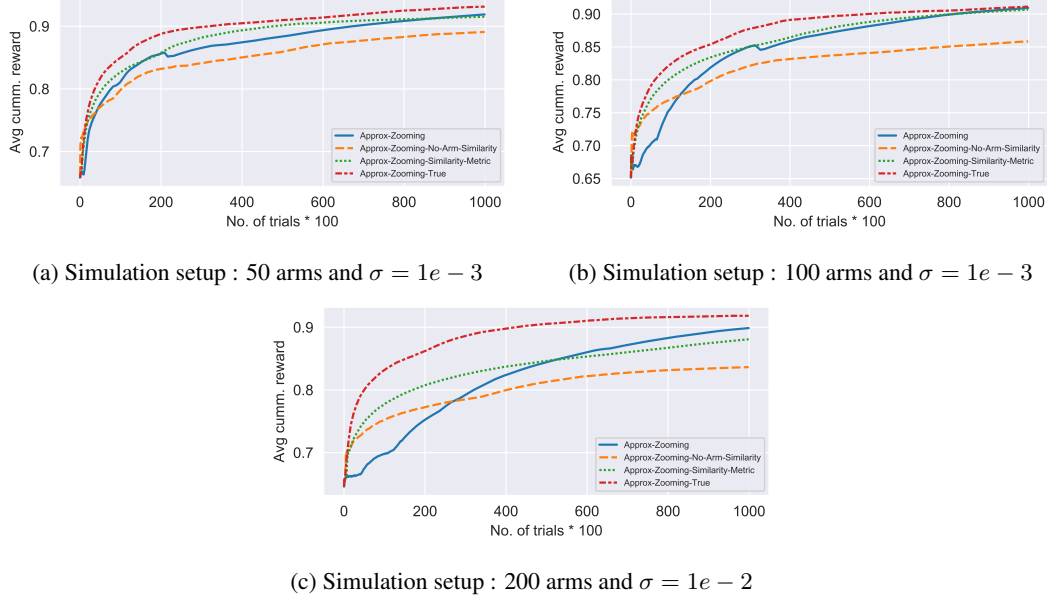




Figure 3: Avg. cumulative reward vs. number of trials

important to use similarities amongst arms to find the optimal arm. Furthermore, we observed the algorithm using the metric representation narrows in slower than the algorithm given true distance function. Therefore, we argue if a metric space is used to find similarities in the arm spaces it needs to be carefully chosen to represent the reward distribution, which is not a trivial task. In contrast, our approach relies on samples from the reward distribution and learns the latent structure avoiding the difficulty of choosing a suitable metric. However, our approach needs to carefully tune the parameter  $k$  to avoid unnecessary sampling for similarity estimation. We anticipate that the benefits of learning the metric only dominates in regimes where the number of arms is large and the time horizon is sufficiently long. Below we include similar plots for other parameters of the problem, in particular we analysed how Approx-Zooming performs compared to other benchmark algorithms when we have smaller number of arms but smaller number of trials with higher values for  $\sigma$ .


In table 1 we plot the final cumulative reward for each of the benchmark algorithms after 10000, 50000, 100000 trials. We have measured the cumulative reward for 50, 100, 200 arms and where  $\sigma$  was set to either  $1e - 1$ ,  $1e - 2$ , or  $1e - 3$ . The hyper parameter values used in each of the benchmark algorithms for this experiment are as follows. For Approx-Zooming algorithm the starting bias was set to 0.25 and the starting distance threshold was set to 0.5. The number of neighbours were set to 10. For the Approx-Zooming-Similarity-Metric algorithm the starting bias was set to 0.25 and the starting distance threshold was set to 10. For the Approx-Zooming-True algorithm the starting bias was set to 0.25 and the starting distance threshold was set to 0.07. As can be seen in table 1 for 50 arms or 100 arms we see a lower cumulative reward for Approx-Zooming compared to other benchmark algorithms when  $\sigma = 1e - 1$  or  $1e - 2$  and the total number of trials is 10,000, which suggests that the cost due to the added extra exploration may exceed the gain from learning the metric when  $\sigma$  is large or the total number of trials is small. This observation is aligned with the click through rates we report in figure 3 and illustrates that the benefits of learning the metric only dominates in regimes where the number of arms is large and the time horizon is sufficiently long.




(a) Approx Zooming



(b) Approx-Zooming-True




(c) Approx-Zooming-Similarity-Metric




(d) Approx-Zooming-No-Arm-Similarity


Figure 4: Arm Frequency Plots For 50 Arms




(a) Approx Zooming



(b) Approx-Zooming-True




(c) Approx-Zooming-Similarity-Metric




(d) Approx-Zooming-No-Arm-Similarity

Figure 5: Arm Frequency Plots For 100 Arms







(a) Approx Zooming



(b) Approx-Zooming-True



(c) Approx-Zooming-Similarity-Metric



(d) Approx-Zooming-No-Arm-Similarity

Figure 6: Arm Frequency Plots For 200 Arms

Number of arms	Total number of trials	$\sigma$	Approx-Zooming	Approx-Zooming-True	Approx-Zooming-Similarity-Metric	Approx-Zooming-No-Arm-Similarity
50 100 200	10000	0.1	0.789319 0.794965 0.663524	0.865815 0.836930 0.838799	0.828892 0.810921 0.788084	0.811413 0.776608 0.757901
50 100 200	50000	0.1	0.905667 0.895329 0.850035	0.911928 0.897861 0.905815	0.900178 0.882766 0.848488	0.861530 0.838076 0.815208
50 100 200	100000	0.1	0.926956 0.917762 0.909856	0.931452 0.911701 0.918141	0.914195 0.906180 0.881553	0.889772 0.859232 0.836007
50 100 200	10000	0.01	0.816736 0.762866 0.735657	0.862454 0.833096 0.838271	0.828892 0.810921 0.788084	0.811956 0.777493 0.760152
50 100 200	50000	0.01	0.893435 0.883395 0.840253	0.911078 0.897688 0.908507	0.903131 0.882611 0.850934	0.861970 0.837934 0.814603
50 100 200	100000	0.01	0.926705 0.912631 0.907565	0.931727 0.911363 0.918062	0.913347 0.906772 0.881263	0.891015 0.858644 0.836181
50 100 200	10000	0.001	0.843526 0.767987 0.703802	0.863722 0.833951 0.837280	0.837772 0.812152 0.787980	0.811940 0.776157 0.760589
50 100 200	50000	0.001	0.892600 0.884865 0.866195	0.910488 0.897461 0.908066	0.902087 0.882370 0.850165	0.864150 0.837778 0.814999
50 100 200	100000	0.001	0.927094 0.908725 0.910485	0.931463 0.911319 0.916554	0.913255 0.905410 0.881567	0.891716 0.859420 0.836579

Table 1: The Cumulative Reward For Different Number Of Arms =  $[50, 100, 200]$ , Trials =  $[10K, 50K, 100K]$  and  $\sigma = [1e-1, 1e-2, 1e-3]$