
Online Algorithm for Unsupervised Sequential Selection with Contextual Information

Arun Verma
Department of IEOR
IIT Bombay, India
v.arun@iitb.ac.in

Manjesh K. Hanawal
Department of IEOR
IIT Bombay, India
mhanawal@iitb.ac.in

Csaba Szepesvári
DeepMind/University of Alberta
Alberta, Canada
szepi@google.com

Venkatesh Saligrama
Department of ECE
Boston University, USA
srv@bu.edu

Abstract

In this paper, we study *Contextual Unsupervised Sequential Selection* (USS), a new variant of the stochastic contextual bandits problem where the loss of an arm cannot be inferred from the observed feedback. In our setup, arms are associated with fixed costs and are ordered, forming a cascade. In each round, a context is presented, and the learner selects the arms sequentially till some depth. The total cost incurred by stopping at an arm is the sum of fixed costs of arms selected and the stochastic loss associated with the arm. The learner's goal is to learn a decision rule that maps contexts to arms with the goal of minimizing the total expected loss. The problem is challenging as we are faced with an unsupervised setting as the total loss cannot be estimated. Clearly, learning is feasible only if the optimal arm can be inferred (explicitly or implicitly) from the problem structure. We observe that learning is still possible when the problem instance satisfies the so-called 'Contextual Weak Dominance' (CWD) property. Under CWD, we propose an algorithm for the contextual USS problem and demonstrate that it has sub-linear regret. Experiments on synthetic and real datasets validate our algorithm.

1 Introduction

Industrial systems, such as those found in medical, airport security, and manufacturing, utilize a suite of tests or classifiers for monitoring patients, people, and products. Tests have costs with the more intrusive and informative ones resulting in higher monetary costs and higher latency. For this reason, they are often organized as a classifier cascade (Chen et al., 2012; Trapeznikov and Saligrama, 2013; Wang et al., 2015), so that new input is first probed by an inexpensive test then a more expensive one. The goal of a cascaded system is to resolve easy to handle examples early so that the overall system maintains high accuracy at low average costs.

Over time, due to environmental changes or test calibrations, sequential testing protocols (STP) may no longer be accurate, resulting in higher costs. While one can leverage off-line methods such as supervised training of cascades (Wang et al., 2015), they require new annotated data collection. In many scenarios, new data cannot be collected in-situ, and system shutdown is not an option. In the absence of annotated data, we face a dilemma. While we can observe test outcomes, we cannot ascertain their reliability due to the absence of ground truth, necessitating *unsupervised sequential selection* (USS) methods, where an arm represents a test/classifier. Recent works (Hanawal et al., 2017; Verma et al., 2019a, 2020a) propose methods for solving the USS problem; however, they

focus exclusively on the non-contextual setting, which in essence requires inputs (people, objects, or products) to be homogeneous, and as such, these methods are unrealistic since contexts (high vs. low risk) can guide the arm selection.

In this context, we propose the contextual USS. In our setup, inputs arrive sequentially, and the learner observes a continuous-valued context as input. While the learner knows the costs of each arm, he does not know the associated stochastic loss. Furthermore, the learner does not benefit from feedback from his arm selection, in contrast to the conventional contextual bandit works (Beygelzimer et al., 2011). Thus, while being agnostic to the true loss, the learner must sequentially choose the arm that leads to the smallest total loss, where the total loss is the sum of the cost of using an arm and the mean loss associated with the arm. As such, our proposed problem is a special case of the stochastic partial monitoring problem with contextual inputs (Lattimore and Szepesvári, 2020, Chapter 37). Most of the prior work on partial monitoring problem is restricted to cases where observed feedback can identify the losses for selected actions. However, in many areas like crowd-sourcing (Bonald and Combes, 2017; Kleindessner and Awasthi, 2018), resource allocation (Verma et al., 2019b), medical diagnosis (Verma et al., 2020b), and many others, feedback from actions may not even be sufficient to identify the losses.

While we draw upon several concepts introduced in earlier work (Hanawal et al., 2017), there are additional challenges in the contextual case due to the unsupervised nature of the problem. First, unlike vanilla-USS, the loss here is context-dependent. We propose notions of contextual weak dominance as a means to relate observed disagreements to differences in losses between any two arms. We then propose a parameterized Generalized Linear Model (GLM) to model the context-conditional disagreement probability between any two arms and validate the model empirically.

A fundamental technical challenge is in the estimation of disagreement probabilities uniformly across all contexts in the finite time while ensuring sufficient exploration between different arm selection protocols, required for honing in on the optimal selection strategy. In particular, since contexts are continuous-valued, and because we have no control over inputs, the contextual observations, in the finite time, may not persistently span the whole space, and estimates are often unreliable. To this end, we adapt techniques from parameterized contextual bandits (Chu et al., 2011; Li et al., 2017) for our unsupervised setting. We propose an algorithm based on the principle of optimism, namely, the larger indexed arm in cascade is chosen when uncertain. We show that our algorithm navigates the exploration-exploitation tradeoffs in different ways and lead to sub-linear cumulative regret. We then validate it on several problem instances derived from synthetic and real datasets.

Related Work. *Stochastic Contextual multi-armed Bandits (SCB)*: In each round, the learner observes the context and decides which arm, among a finite number of arms, to apply (Beygelzimer et al., 2011). By playing an arm, the learner observes a stochastic reward that depends on the context and the arm selected. The most commonly studied model assumes that each arm is parameterized, and the mean reward of an arm is the inner product of the context and an unknown parameter associated with the arm. Contextual bandits have been applied to problems ranging from online advertising (Li et al., 2010; Chu et al., 2011) and recommendations (Langford and Zhang, 2008) to clinical trials (Woodroffe, 1979) and mobile health (Tewari and Murphy, 2017). *Generalized linear models (GLM)* assume that the mean reward is a non-linear link function of the inner product between the context vector and the unknown parameter vector (Filippi et al., 2010; Li et al., 2017). GLMs are also useful models for the classification problems where rewards, in the context of online learning problems, could be binary (Zhang et al., 2016; Jun et al., 2017). A more challenging non-parameterized version of the stochastic contextual bandits is studied in (Agarwal et al., 2014).

Another framework that is closely related to SCB is *stochastic linear bandits (SLB)* (Auer, 2002; Dani et al., 2008; Rusmevichientong and Tsitsiklis, 2010; Abbasi-Yadkori et al., 2011). In this setup, the environment is parameterized, and there could be uncountably many arms (within some bounded radius), also referred to as decision set. The arms are characterized into their feature vectors, and the mean reward for playing an arm is given as the inner product of the parameter (unknown) and the feature vector associated with the arm. In situations where the decision set is allowed to vary in each round and are finite, SLBs are equivalent to SCBs, where feature vectors correspond to context-arm pairs (Li et al., 2010, 2017). For our work, we leverage GLMs as models for disagreement probability between any two arms. While it is tempting to reduce contextual USS to SCBs, note that, unlike prior works, we do not observe loss for our action choices, and so conventional algorithms such as LinUCB and UCB-GLM (Li et al., 2010; Agarwal et al., 2014; Li et al., 2017) cannot be applied.

Most of the prior work (Hanawal et al., 2017; Verma et al., 2019a, 2020a) considered the problem of learning an optimal action but ignored the contextual information. In this work, we incorporated contextual information, which is readily available in many applications. Exploiting the *real-valued* contextual information (features) for improving the arm selection strategy is non-trivial due to the unsupervised nature of the problem where the standard analysis of contextual bandits does not apply. We made necessary modeling assumptions to leverage GLMs to parameterize the disagreement probability between two arms and extended the existing definitions to address the new setup’s learnability issues. However, the problem still requires new ideas and analysis methods to derive an efficient algorithm, which poses new technical challenges for analysis.

2 Problem Setting

We consider a stochastic contextual bandits problem with K arms. The set of arms is denoted as $[K]$ where $[K] \doteq \{1, 2, \dots, K\}$. In each round t , the environment generates a vector $(X_t, Y_t, \{Y_t^i\}_{i \in [K]})$. The vector X_t denotes the context in round t and forms an independent and identically distributed (IID) sequence drawn from a bounded set $\mathcal{X} \subset \mathbb{R}^d$ according to an unknown but fixed distribution ν . The binary reward for context X_t is denoted by $Y_t \in \{0, 1\}$, which is hidden from the learner. The vector $(\{Y_t^i\}_{i \in [K]}) \in \{0, 1\}^K$ represents observed feedback at time t , where Y_t^i denotes the feedback observed after playing arm i with X_t as input¹. We denote the cost for using arm i as $c_i \geq 0$ that is known and the same for all contexts.

In contextual USS, the arms are assumed to be ordered and form a cascade. When the learner selects an arm $i \in [K]$, the feedback from all arms till arm i in the cascade are observed. The expected loss of playing the arm i for a given context x_t is denoted as $\gamma_i(x_t) \doteq \mathbb{E} \left[\mathbb{1}_{\{Y_t^i \neq Y_t | X = x_t\}} \right] = \mathbb{P} \{Y_t^i \neq Y_t | X = x_t\}$, where $\mathbb{1}_{\{A\}}$ denotes indicator of event A . For soundness, we assume that the probability density function of context distribution is strictly positive on $\mathcal{X}$ such that the conditional probabilities are well defined. The total expected loss incurred by playing arm i for context x_t is defined as $\gamma_i(x_t) + \lambda_i C_i$, where $C_i \doteq c_1 + \dots + c_i$ and λ_i is a trade-off parameter that normalizes the incurred cost and the loss of playing arm i .

Since the true rewards are hidden from the learner, the expected loss of an arm cannot be inferred from the observed feedback. We thus have a version of the stochastic partial monitoring problem (Cesa-Bianchi et al. (2006); Bartók and Szepesvári (2012); Bartók et al. (2014)), and we refer to it as contextual unsupervised sequential selection (USS). Let $\mathbf{Q}$ be the unknown joint distribution of $(X, Y, Y^1, Y^2, \dots, Y^K)$. Henceforth we identify a contextual USS instance as $P \doteq (\mathbf{Q}, \mathbf{c})$ where $\mathbf{c} \doteq (c_1, c_2, \dots, c_K)$ is the known cost vector of arms. We denote the collection of contextual USS instances as $\mathcal{P}_{\text{USS}}$. For instance $P \in \mathcal{P}_{\text{USS}}$, the optimal arm for a context x_t is given as follows:

$$i_t^* \in \max \left\{ \arg \min_{i \in [K]} (\gamma_i(x_t) + \lambda_i C_i) \right\}, \quad (1)$$

where the choice of i_t^* is risk-averse as we prefer the arm with lower error among the optimal arms.

The interaction between the environment and a learner is given in Algorithm 1.

Algorithm 1 Learning on contextual USS instance $(\mathbf{Q}, \mathbf{c})$

For each round t :

1. **Environment** chooses a vector $(X_t, Y_t, \{Y_t^i\}_{i \in [K]}) \sim \mathbf{Q}$.
 2. **Learner** observes a context $X_t = x_t$ and selects an arm $I_t \in [K]$ to stop in cascade.
 3. **Feedback and Loss:** The learner observes feedback $(Y_t^1, Y_t^2, \dots, Y_t^{I_t})$ and incurs a total loss $\mathbb{1}_{\{Y_t^{I_t} \neq Y_t | X = x_t\}} + \lambda_{I_t} C_{I_t}$.
-

¹In our setup, an arm i could be a classifier that outputs label Y^i . The classifier’s input could be a context and any combinations of feedback observed from classifiers coming before the arm i in the cascade. For example, consider a case where each arm represents a crowd-sourced worker. After using the first i crowd-sourced workers, the final label can be a function of predicted labels of the first i crowd-sourced workers.

The learner's goal is to find an arm for each context such that the cumulative expected loss is minimized. Specifically, for T contexts, we measure the performance of a policy that selects an arm I_t for a context x_t in terms of regret given by

$$\mathfrak{R}_T = \sum_{t=1}^T (\gamma_{I_t}(x_t) + \lambda_{I_t} C_{I_t} - (\gamma_{i_t^*}(x_t) + \lambda_{i_t^*} C_{i_t^*})). \quad (2)$$

We seek policies that yield sub-linear regret, i.e., $\mathfrak{R}_T/T \rightarrow 0$ as $T \rightarrow \infty$. It implies that the learner collects almost as much reward in the long run as an oracle collects that knew the optimal arm for every context. We say that a problem instance $P \in \mathcal{P}_{\text{USS}}$ is learnable if there exists a policy such that $\lim_{T \rightarrow \infty} \mathfrak{R}_T/T = 0$.

In the sequel, we discuss the selection criteria for optimal arm for a given context and the conditions under which instances of $\mathcal{P}_{\text{USS}}$ are learnable.

2.1 Contextual Weak Dominance

Next, we introduce the contextual weak dominance property of a problem instance.

Definition 1 (Contextual Weak Dominance (CWD)). *Let i_t^* denote optimal arm for context x_t . Then the context x_t is said to satisfy weak dominance (WD) property if*

$$\forall j > i_t^* : C_j - C_{i_t^*} > \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\}. \quad (3)$$

A problem instance $P \in \mathcal{P}_{\text{USS}}$ is said to satisfy the CWD property if all contexts of P satisfy WD property. We denote the set of all instances in $\mathcal{P}_{\text{USS}}$ that satisfies CWD property by $\mathcal{P}_{\text{CWD}}$.

In the following, we use an alternative characterization of the CWD property, given as

$$\xi(x_t) \doteq \min_{j > i_t^*} \left\{ C_j - C_{i_t^*} - \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\} \right\} > 0. \quad (4)$$

We define $\xi \doteq \inf_{x \in \mathcal{X}} \xi(x)$ and assume that $\xi > 0$. The larger the value of ξ , 'stronger' is the CWD property, and easier it is to identify an optimal arm for given contexts. We later characterize the regret upper bounds of proposed algorithms in terms of ξ . We also discuss the case when a fraction of contexts satisfies WD property in the supplementary material.

2.2 Selection Criteria for Optimal Arm

Without loss of generality, we set $\lambda_i = 1$ for all $i \in [K]$ as their value can be absorbed into the costs. Since $i_t^* = \arg \min_{i \in [K]} (\gamma_i(x_t) + C_i)$, it must satisfy following equation:

$$\forall j < i_t^* : C_{i_t^*} - C_j \leq \gamma_j(x_t) - \gamma_{i_t^*}(x_t), \quad (5a)$$

$$\forall j > i_t^* : C_j - C_{i_t^*} > \gamma_{i_t^*}(x_t) - \gamma_j(x_t). \quad (5b)$$

As the loss of an arm is not observed, the above equations can not lead to a sound arm selection criteria. We thus have to relate the unobservable quantities in terms of the quantities that can be observed. In our setup, we can compare the feedback of two arms, which can be used to estimate the disagreement probabilities between them. For notation convenience, we define $p_{ij}^{(t)} \doteq \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$ for $i < j$. The value of $p_{ij}^{(t)}$ can be estimated as it is observable. Our next result bounds unobserved error rates differences in terms of their observable disagreement probabilities for a given context.

Lemma 1. *For any i, j , and $x_t \in \mathcal{X}$, $\gamma_i(x_t) - \gamma_j(x_t) = p_{ij}^{(t)} - 2\mathbb{P} \left\{ Y_t^i = Y_t, Y_t^j \neq Y_t | X = x_t \right\}$.*

The detailed proof of Lemma 1 and all other missing proofs appear in the supplementary material.

Now, using Lemma 1, we can replace Eq. (5a) by

$$\forall j < i_t^* : C_{i_t^*} - C_j \leq p_{ji_t^*}^{(t)}, \quad (6)$$

which only has observable quantities. For $j > i_t^*$, using the CWD property, we replace Eq. (5b) by

$$\forall j > i_t^* : C_j - C_{i_t^*} > p_{i_t^* j}^{(t)}. \quad (7)$$

Using Eq. (6) and Eq. (7), our next result gives the optimal arm for a given context x_t .

Lemma 2. *Let $P \in \mathcal{P}_{\text{CWD}}$ and $\mathcal{B}_t = \left\{ i : \forall j > i, C_j - C_i > p_{ij}^{(t)} \right\} \cup \{K\}$. Then the arm $I_t = \min(\mathcal{B}_t)$ is the optimal arm for a context x_t .*

By construction, the optimal arm lies in set $\mathcal{B}_t$. Because of Eq. (6), any sub-optimal arm having smaller index than optimal arm do not satisfy Eq. (7), hence it can not be in set $\mathcal{B}_t$. Therefore, the smallest arm of set $\mathcal{B}_t$ is the optimal arm.

Theorem 1. *The set $\mathcal{P}_{\text{CWD}}$ is maximal learnable.*

The proof establishes that under the CWD property, there exists a ‘sound’ arm selection policy that identifies the optimal arm for each context. The sound policy only uses conditional disagreement probabilities between pairs of arms that can be estimated from the feedback of arms.

3 Parameterization of Pairwise Disagreement Probability

Since the number of contexts could be much larger (can be infinite) than the learning horizon, in stochastic contextual bandits, a correlation structure is assumed between the reward (loss) and the contexts (Auer, 2002; Li et al., 2010, 2017). It is often realized via parameterization of the arms such that expected rewards (or losses) observed from an arm depend on the unknown parameter. In our setting, we cannot observe a loss for any arm. Hence parameterization of an expected loss of the arms is not useful. However, we can obtain feedback of two arms for a given context and can compare them. For example, we can check whether two arms’ feedback agrees or disagrees for a given context. Thus, we assume a correlation structure on the disagreement probability for a pair of arms across the contexts and parameterize it using generalized linear models. For $i < j$ and context x_t , the disagreement probability for (i, j) pair of arms is given via a function μ as follows:

$$\mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\} = \mu(\Phi_{ij}(x_t)^\top \theta_{ij}^*), \quad (8)$$

where $x_t \in \mathbb{R}^d$, $\Phi_{ij} : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ is a feature map for some $d' \geq d$,² and $\theta_{ij}^* \in \mathbb{R}^{d'}$ is the unknown parameter for (i, j) pair.

We assume the following assumptions on context distribution ν and function μ , which is standard in the GLM bandit literature (Filippi et al., 2010; Li et al., 2017):

Assumption 1 (GLM). • For all $x \in \mathcal{X}$ and (i, j) pairs, $\|\Phi_{ij}(x)\|_2 \leq 1$.

- $\kappa \doteq \inf_{\|x\|_2 \leq 1, \|\theta - \theta_{ij}^*\|_2 \leq 1} \dot{\mu}(\Phi_{ij}(x)^\top \theta) > 0$ for all (i, j) pairs.
- There exists a constant $\lambda_\Sigma > 0$ such that $\lambda_{\min}(\mathbb{E}[\Phi_{ij}(X)\Phi_{ij}(X)^\top]) \geq \lambda_\Sigma$ for all (i, j) pairs.
- The function $\mu : \mathbb{R} \rightarrow [0, 1]$ is continuously differentiable and Lipschitz with constant k_μ .

For our setting, the function μ is defined as $\mu(z) = 1/(1 + e^{-z})$, which is the logistic function. The logistic function is widely used function for binary classification model and has $k_\mu \leq 1/4$.

In contextual USS setup, we can compare the arms’ feedback and check whether they agree or not for a given context. These binary observations (agree or disagree) can be treated as noisy samples of the disagreement probability. The noise in the binary observation obtained by comparing the feedback of (i, j) pair of arms in round t , is given by

$$\varepsilon_{ij}^{(t)} = \begin{cases} 1 - \mu(\Phi_{ij}(x_t)^\top \theta_{ij}^*), & \text{with probability } \mu(\Phi_{ij}(x_t)^\top \theta_{ij}^*) \\ -\mu(\Phi_{ij}(x_t)^\top \theta_{ij}^*), & \text{with probability } (1 - \mu(\Phi_{ij}(x_t)^\top \theta_{ij}^*)) \end{cases}$$

²Let $\mathbb{R}^{d_{ij}}$ be the space where Eq. (8) holds for (i, j) pair of arms, and Φ_{ij} is the feature map that lift x_t from $\mathbb{R}^d$ space to $\mathbb{R}^{d_{ij}}$ space. For simplicity, we take $d' = \max_{i < j \leq K} d_{ij}$.

where $\varepsilon_{ij}^{(t)}$ is $\mathcal{F}_t$ -measurable with $\mathbb{E}[\varepsilon_{ij}^{(t)}|\mathcal{F}_t] = 0$. Here $\mathcal{F}_t$ denotes sigma algebra generated by history $\left\{ \left(X_s, I_s, \{Y_s^i\}_{i \in [I_s]} \right) \right\}_{s \in [t]}$ till time t . Since $\varepsilon_{ij}^{(t)}$ is a zero-mean shifted Bernoulli random variable, $\varepsilon_{ij}^{(t)}$ satisfies the following sub-Gaussian condition with parameter $\sigma \in (0, 1)$:

$$\mathbb{E} \left[\exp(\lambda \varepsilon_{ij}^{(t)}) | \mathcal{F}_t \right] \leq \exp \left(\frac{\lambda^2 \sigma^2}{2} \right), \quad \forall \lambda \in \mathbb{R}.$$

Let $d_{ij}(t) \doteq \mathbb{1}_{\{Y_t^i \neq Y_t^j | X = x_t\}}$ be the disagreement indicator for a context x_t and S_{ij}^t be the set of indices of contexts for which disagreements are observed for (i, j) pair of arms till round t . In round t , we estimate θ_{ij}^* , denoted by $\hat{\theta}_{ij}^t$, using the following equation adapted from the maximum likelihood estimator (MLE) used for GLM bandits (Filippi et al., 2010; Li et al., 2017):

$$\sum_{s \in S_{ij}^t} (d_{ij}(s) - \mu(\Phi_{ij}(x_s)^\top \theta)) \Phi_{ij}(x_s) = 0. \quad (9)$$

In the next section, we develop an algorithm that exploits Lemma 2 for selecting the optimal arm to each context. The algorithm replaces the terms $p_{ij}^{(t)}$ in Lemma 2 by their optimistic estimates.

4 Algorithm for Contextual USS: USS-PD

Our algorithm, named **USS-PD**, is based on the *optimism-in-the-face-of-uncertainty* (OFU) principle. **USS-PD** works as follows: It takes δ and m as inputs, where δ is the confidence in the estimated parameters and used for computing confidence bound for θ_{ij}^* as given by Lemma 4. The choice of m ensures that with probability at least $(1 - \delta)$, the sample correlation matrix $V_{ij}^t = \sum_{s \in S_{ij}^t} \Phi_{ij}(x_s) \Phi_{ij}(x_s)^\top$ for each (i, j) pair where $i < j$, is invertible. A high probability upper bound on m is computed using Lemma 3. The algorithm collects feedback from all arms by selecting the arm K irrespective of the context received for first m rounds. After m rounds, the sample correlation matrix and the estimate of θ_{ij}^* are computed for each (i, j) pair where $i < j$.

For $t > m$, the learner receives a context x_t and plays the arm $i = 1$ and then observe its feedback. For each (i, j) pair and context x_t , the upper bound on disagreement probability $\tilde{p}_{ij}^{(t)}$ is computed using $\hat{\theta}_{ij}^t$ and confidence bonus $\alpha_{ij}^t \|\Phi_{ij}(x_t)\|_{(V_{ij}^t)^{-1}}$. Here the notation $\|x\|_A^2 \doteq x^\top A x$ denotes the weighted l_2 -norm of vector $x \in \mathbb{R}^d$ with respect to a positive definite matrix $A \in \mathbb{R}^{d \times d}$. The confidence bonus has two terms. The first term α_{ij}^t is a slowly increasing function in t whose value is specified in Lemma 4, and the second term $\|\Phi_{ij}(x_t)\|_{(V_{ij}^t)^{-1}}$ decreases to zero as t increases.

USS-PD Algorithm for Contextual USS using Pairwise Disagreement

- 1: **Input:** Tuning parameters: $\delta \in (0, 1)$ and $m > 0$
 - 2: Select arm K for first m contexts
 - 3: $\forall i < j \leq K$: set $V_{ij}^m \leftarrow \sum_{t=1}^m \Phi_{ij}(x_t) \Phi_{ij}(x_t)^\top$ and update $\hat{\theta}_{ij}^m$ by solving Eq. (9)
 - 4: **for** $t = m + 1, m + 2, \dots$ **do**
 - 5: Receive context x_t . Set $i = 1$ and $I_t = 0$
 - 6: **do**
 - 7: Play arm i
 - 8: $\forall j \in [i + 1, K]$: compute $\tilde{p}_{ij}^{(t)} \leftarrow \mu \left(\Phi_{ij}(x_t)^\top \hat{\theta}_{ij}^{t-1} + \alpha_{ij}^{t-1} \|\Phi_{ij}(x_t)\|_{(V_{ij}^{t-1})^{-1}} \right)$
 - 9: If $\forall j \in [i + 1, K]$: $C_j - C_i > \tilde{p}_{ij}^{(t)}$ or $i = K$ then set $I_t = i$ else set $i = i + 1$
 - 10: **while** $I_t = 0$
 - 11: Select arm I_t and observe $Y_t^1, Y_t^2, \dots, Y_t^{I_t}$
 - 12: $\forall i < j \leq I_t$: update $V_{ij}^t \leftarrow V_{ij}^{t-1} + \Phi_{ij}(x_t) \Phi_{ij}(x_t)^\top$ and $\hat{\theta}_{ij}^t$ by solving Eq. (9)
 - 13: **end for**
-

After computing $\tilde{p}_{ij}^{(t)}$, the algorithm checks whether the arm i is the best arm using Eq. (7) with $\tilde{p}_{ij}^{(t)}$ in place of $p_{ij}^{(t)}$. If the arm i is not the best, then the algorithm plays the next arm, and then the same process is repeated. If the arm i is the best arm for context, then the algorithm stops at that arm with $I_t = i$ for that context. After selecting arm I_t , the feedback from arms $1, \dots, I_t$ are observed. After that, the values of V_{ij}^t are updated, and $\hat{\theta}_{ij}^t$ are re-estimated. The same process is repeated for subsequent contexts.

Remark 1. GLM bandits are well studied but require reward or loss information. In the USS setup, loss of selected arm can not be observed; hence finding the optimal arm is challenging. Due to binary disagreement, **USS-PD** uses the MLE estimator for θ_{ij}^* as used in GLM bandits (Filippi et al., 2010; Li et al., 2017). However, the feedback structure and the way arms are selected in the USS setup differ from that in the GLM bandits. Further, our analysis needs carefully connecting the regret with the bad events that make **USS-PD** selects non-optimal arms.

Remark 2. We force the algorithm to explore until the correlation matrix V_{ij}^t is invertible for all (i, j) pairs. The invertibility can also be ensured by adding a regularization term (Abbasi-Yadkori et al., 2011; Zhang et al., 2016; Jun et al., 2017) to avoid forced exploration. However, the analysis of **USS-PD** with regularization term still required to the non-regularized part of the sample correlation matrix becomes invertible. See the supplementary material for the algorithm and its analysis.

4.1 Regret Analysis of **USS-PD**

The following definition is useful in our regret analysis.

Definition 2 (Arm Preference ($\succ_t$)). **USS-PD** prefers an arm i over j for context x_t if

$$i \succ_t j \doteq \begin{cases} C_i - C_j < \tilde{p}_{ji}^{(t)}, & \text{if } j < i \\ C_j - C_i > \tilde{p}_{ij}^{(t)}, & \text{if } j > i. \end{cases} \quad (10a)$$

$$(10b)$$

Our next result bounds the number of disagreement observations required from a pair of arms say (i, j) , such that the smallest eigenvalue of its sample correlation matrix V_{ij} matrices is larger than a fixed value. This result uses the standard results from random matrix theory (Vershynin, 2012).

Lemma 3. Let $V_{ij}^t = \sum_{s \in S_{ij}^t} \Phi_{ij}(x_s) \Phi_{ij}(x_s)^\top$, $\Sigma_{ij} = \mathbb{E} [\Phi_{ij}(X) \Phi_{ij}(X)^\top]$, Ψ and $\delta \in (0, 1)$ be two positive constants. Then, there exist positive universal constants C_1 and C_2 such that the minimum eigenvalue of $\lambda_{\min}(V_{ij}^t) \geq \Psi$ with probability at least $1 - 2\delta/K^2$, iff

$$|S_{ij}^t| \geq \left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log(K^2/2\delta)}}{\lambda_{\min}(\Sigma_{ij})} \right)^2 + \frac{2\Psi}{\lambda_{\min}(\Sigma_{ij})}.$$

The next result is adapted to our setting from the confidence bounds for maximum likelihood estimator used in GLM bandits (Li et al., 2017).

Lemma 4 (Confidence Ellipsoid). Let m be such that $\lambda_{\min}(V_{ij}^{m+1}) \geq 1$ for any pair (i, j) . Then the following event holds with probability at least $1 - 2\delta/K^2$ for **USS-PD**:

$$\left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{V_{ij}^t} \leq \alpha_{ij}^t, \quad \forall t > m$$

where $\alpha_{ij}^t = \frac{2\sigma}{\kappa} \sqrt{\frac{d'}{2} \log\left(1 + \frac{2t}{d'}\right) + \log\left(\frac{K^2}{2\delta}\right)}$.

The regret analysis of GLM bandits hinges on bounding the instantaneous regret in each round, which is tied to the estimation error of the GLM parameters. Due to the unsupervised setting and cascade structure, this way of bounding regret does not work in our setup. Our analysis goes by bounding the number of pulls of the sub-optimal arms. However, unlike standard bandits, we have to distinguish whether the sub-optimal arm pulled by **USS-PD** is on the ‘left’ or ‘right’ of the optimal arm in the cascade. It requires our analysis to handle both the cases carefully. Since **USS-PD** uses a similar MLE estimator for parameter estimation as in GLM bandits, we only adapt their asymptotic normality results. Our next results give conditions when **USS-PD** prefers a sub-optimal arm for a context.

Lemma 5. Let $\theta \in \Theta_{\text{CWD}}$. Then **USS-PD** prefers any sub-optimal arm $l < i_t^*$ for context x_t with probability at most $\delta/2$.

Lemma 6. Let $\theta \in \Theta_{\text{CWD}}$. If **USS-PD** prefers a sub-optimal arm $h > i_t^*$ for context x_t then

$$2k_\mu \alpha_{i_t^* h}^t > \xi_{i_t^* h}(x_t) \sqrt{\lambda_{\min}(V_{i_t^* h}^t)}.$$

where $\xi_{i_t^* h} = C_h - C_{i_t^*} - p_{i_t^* h}^{(t)}$ and α_{ij}^t is given by Lemma 4.

Let $m \doteq C\lambda_\Sigma^{-2} (d' + \log(k^2/2\delta)) + 2\lambda_\Sigma^{-1}$, where $C > 0$ is the universal constant and $R_{\max} \doteq \max_{i \in [K], x \in \mathcal{X}} [C_i + \gamma_i(x) - (C_{i^*} + \gamma_{i^*}(x))]$, where i^* is the optimal arm for context x . Now we state the regret upper bound of **USS-PD**.

Theorem 2 (Regret Upper Bound). Let $\theta \in \Theta_{\text{CWD}}$, $\delta \in (0, 1)$, Assumption 1 holds, and $\xi_h = \min_{t \geq 1} \xi_{i_t^* h}(x_t)$. Then with probability at least $1 - 2\delta$, the regret of **USS-PD** for $T > m$ contexts is

$$\begin{aligned} \mathfrak{R}_T \leq R_{\max} & \left[m + \sum_{h=2}^K \left(\left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log\left(\frac{K^2}{2\delta}\right)}}{\lambda_\Sigma} \right)^2 + \frac{16}{\lambda_\Sigma} \right. \right. \\ & \left. \left. \left(\frac{k_\mu \sigma}{\xi_h \kappa} \right)^2 \left(\frac{d'}{2} \log \left(1 + \frac{2T}{d'} \right) + \log \left(\frac{K^2}{2\delta} \right) \right) \right] \right]. \end{aligned}$$

Corollary 1. Let technical conditions stated in Theorem 2 hold. Then with probability at least $1 - 2\delta$

$$\mathfrak{R}_T \leq O(Kd' \log(T)/\xi^2).$$

The regret of **USS-PD** for instance $\theta \in \Theta_{\text{CWD}}$ is logarithmic in T and grows linearly with d' and K . The regret is inversely dependent on the value of $\xi \doteq \min_{h \geq 2} \xi_h$ (measure how well CWD holds),

which implies the problem instance with smaller ξ has more regret and vice-versa. The value of ξ is analogous to the minimum sub-optimality gap in the standard Multi-Armed Bandits setting. With a large context set, ξ can be small, and its inverse relation in the regret captures the difficulty of the USS problem.

5 Experiment

We evaluate the performance of **USS-PD** on different problem instances derived from synthetic and real datasets. In our experiments, the data samples are treated as contexts. The labels of contexts are known but are never revealed to the algorithm. We use the labels to train classifiers offline that act as arms. Arm i represents a logistic classifier with trained parameter θ_i . A context (data sample) x is assigned label 1 from the i -th classifier with probability $\mu(x^\top \theta_i)$ and label 0 with probability $1 - \mu(x^\top \theta_i)$. The disagreement labels for (i, j) pair is computed using the labels of classifier i and j . To satisfy Eq. (8), we use the polynomial kernel of degree two for mapping context into higher-dimensional space. Unlike other kernels, the polynomial kernel uses a well-defined feature map to lift the contexts into fixed, higher-dimensional space. The details of the used problem instances are as follows.

Synthetic Dataset: We consider 3-dimensional synthetic dataset with 5000 data samples. Each sample is represented by $x = (x_1, x_2, x_3)$, where the value of x_j is drawn uniformly at random from $(-1, 1)$. A sample x is labeled 0 if the value of $(x_1 + x_1 x_2 + x_3^2)$ is negative otherwise it is labeled 1. We train five logistic classifiers on this synthetic dataset by varying the regularization parameter. We then assign a positive cost to each classifier and order them by their increasing cost. We vary the cost of using classifiers to get different problem instances (see details in the supplementary material).

Real Datasets: We applied our algorithm on PIMA Indian Diabetes (Kaggle, 2016) dataset. Each sample has 8 features related to the conditions of the patient. We split the features into three subsets and train a logistic classifier on each subset. We associate 1st classifier with the first 6 features as input. These features include patient history/profile. The 2nd classifier, in addition to the 6 features, utilizes the feature on the glucose tolerance test, and the 3rd classifier uses all the previous features and the feature that gives values of insulin test. Due to space constraints, the experiment results on Heart Disease dataset (Detrano, 1998; Dheeru and Karra Taniskidou, 2017) are given in the supplementary material.

5.1 Experiments Results

We compare the performance of **USS-PD** on four problem instances derived from the synthetic dataset. The instances vary based on the cost of arms. All contexts in Instance 1 do not satisfy WD property; hence it suffers linear regret as shown in Fig. 1a. For the remaining instances, we set costs such that the value of ξ increasing from Instance 2 to 4. As expected, the regret decreases from Instance 2 to 4, as seen in Fig. 1a. We also compare **USS-PD** against an algorithm where the learner receives true labels as feedback. In particular, the learner knows whether the classifier’s output is correct or not and can estimate their error rates. We implement this ‘supervised’ setting by replacing disagreement probability in Eq. (7) with estimated error rates. As expected, the regret with supervision has lower than the **USS-PD** regret (unsupervised) in Fig. 1b. It is qualitatively interesting because these plots demonstrate that, in typical cases, our unsupervised algorithm can eventually learn to perform as good as an algorithm with knowledge of true labels.




Figure 1: Performance of **USS-PD** on different problem instances derived from synthetic and real datasets.

We derive three problem instances from PIMA Indian Diabetes dataset by varying the costs of using classifiers. Since all contexts of these problem instances do not satisfy WD property (see details in the supplementary material), we used cumulative total expected cost as a performance measure, where the cumulative total expected cost is given by $\sum_{t=1}^T (\gamma_{I_t}(x_t) + C_{I_t})$. We compare the performance of **USS-PD** with three baseline policies – the first baseline policy uses the third classifier irrespective of contexts, and it is denoted as policy ‘ $I_t = 3$ ’ (plays arm 3 in each round). The second baseline policy uses the second classifier for all contexts, and it is denoted as policy ‘ $I_t = 2$ ’. The third baseline policy is ‘Random,’ which selects an arm uniformly at random in each round. In all three problem instances, we observe that **USS-PD** performs better than the baselines, as shown in Fig. 1c.

We repeat each of the above experiments 100 times, and then the average regret is presented with a 95% confidence interval. The vertical line on each plot shows the confidence interval.

6 Conclusion and Future Directions

We studied the unsupervised sequential selection problem with contextual information. It is a partial monitoring stochastic contextual bandit problem, where the loss of an arm can not be inferred from the observed feedback. But one can compare the feedback of two arms to see if they agree or disagree. We modeled the disagreement probability between each pair of the arms as linearly parameterized and developed an algorithm named **USS-PD** that achieves $O(\log T)$ regret with high probability.

We exploited the contextual information but ignored the inherent side observations due to the arms’ cascade structure. By using the side observations, one can tighten the regret bounds. Another interesting future direction is to develop algorithms that decide whether it needs to go further down in the cascade when more information about context is revealed along the cascade.

7 Broader Impact

The work considered the unsupervised sequential selection problem with contextual information. While we are not targeting any specific applications, the work has many potential civilian applications.

As usual, these can improve societal conditions, but of course, with any technology, specific deployments need care. However, this is outside of the scope of the present work, which is aimed at improving the basic algorithms and understand the fundamental challenges in this problem setting. Of course, the authors hope that their work will have an altogether positive impact, both by deepening our understanding of challenging sequential decision making under uncertainty and by potential future (careful) applications of the algorithms developed here. Having said this, we do not foresee any immediate negative impact of this work.

Acknowledgments

Manjesh K. Hanawal would like to thank the support from INSPIRE faculty fellowships from DST, Government of India, SEED grant (16IRCCSG010) from IIT Bombay, and Early Career Research (ECR) Award from SERB. Csaba Szepesvári gratefully acknowledges funding from the Canada CIFAR AI Chairs Program, Amii, and NSERC. Venkatesh Saligrama would like to acknowledge NSF Grants DMS -2007350 (VS), CCF-2022446, CCF-1955981, and the Data Science Faculty Fellowship from the Rafik B. Hariri Institute.

References

- M. Chen, Z. Xu, K. Q. Weinberger, O. Chapelle, and D. Kiedem. Classifier cascade: Tradeoff between accuracy and feature evaluation cost. In *International Conference on Artificial Intelligence and Statistics*, pages 235–242, 2012.
- Kirill Trapeznikov and Venkatesh Saligrama. Supervised sequential classification under budget constraints. In *Artificial Intelligence and Statistics*, pages 581–589, 2013.
- Joseph Wang, Kirill Trapeznikov, and Venkatesh Saligrama. Efficient learning by directed acyclic graph for resource constrained prediction. In *Advances in Neural Information Processing Systems* 28, pages 2152–2160. 2015.
- Manjesh Hanawal, Csaba Szepesvari, and Venkatesh Saligrama. Unsupervised sequential sensor acquisition. In *Artificial Intelligence and Statistics*, pages 803–811, 2017.
- Arun Verma, Manjesh K Hanawal, Csaba Szepesvari, and Venkatesh Saligrama. Online algorithm for unsupervised sensor selection. In *Artificial Intelligence and Statistics*, pages 3168–3176, 2019a.
- Arun Verma, Manjesh K Hanawal, and Nandyala Hemachandra. Thompson sampling for unsupervised sequential selection. In *Asian Conference on Machine Learning*, pages 545–560. PMLR, 2020a.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert E. Schapire. Contextual bandit algorithms with supervised learning guarantees. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Thomas Bonald and Richard Combes. A minimax optimal algorithm for crowdsourcing. In *Advances in Neural Information Processing Systems*, pages 4352–4360, 2017.
- Matthäus Kleindessner and Pranjal Awasthi. Crowdsourcing with arbitrary adversaries. In *International Conference on Machine Learning*, pages 2713–2722, 2018.
- Arun Verma, Manjesh Hanawal, Arun Rajkumar, and Raman Sankaran. Censored semi-bandits: A framework for resource allocation with censored feedback. In *Advances in Neural Information Processing Systems*, pages 14499–14509, 2019b.
- Arun Verma, Manjesh K Hanawal, and Nandyala Hemachandra. Unsupervised online feature selection for cost-sensitive medical diagnosis. In *2020 International Conference on COMMunication Systems & NETWORKS (COMSNETS)*, pages 1–6. IEEE, 2020b.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert E Schapire. Contextual bandits with linear payoff functions. In *International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.

- Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning*, pages 2071–2080, 2017.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670. ACM, 2010.
- John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. In *Advances in neural information processing systems*, pages 817–824, 2008.
- Michael Woodroofe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.
- Ambuj Tewari and Susan A Murphy. From ads to interventions: Contextual bandits in mobile health. In *Mobile Health*, pages 495–517. Springer, 2017.
- Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems*, pages 586–594, 2010.
- Lijun Zhang, Tianbao Yang, Rong Jin, Yichi Xiao, and Zhi-hua Zhou. Online stochastic linear optimization under one-bit feedback. In *International Conference on Machine Learning*, pages 392–401, 2016.
- Kwang-Sung Jun, Aniruddha Bhargava, Robert Nowak, and Rebecca Willett. Scalable generalized linear bandits: Online computation and hashing. In *Advances in Neural Information Processing Systems*, pages 99–109, 2017.
- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert E. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, 2014.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *COLT*, pages 355–366, 2008.
- Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010.
- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- Nicolo Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Regret minimization under partial monitoring. *Mathematics of Operations Research*, 31(3):562–580, 2006.
- Gábor Bartók and Csaba Szepesvári. Partial monitoring with side information. In *International Conference on Algorithmic Learning Theory*, pages 305–319. Springer, 2012.
- Gábor Bartók, Dean P Foster, Dávid Pál, Alexander Rakhlin, and Csaba Szepesvári. Partial monitoring—classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 39(4):967–997, 2014.
- Roman Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, page 210–268. Cambridge University Press, 2012. doi: 10.1017/CBO9780511794308.006.
- UCI Machine Learning, Kaggle. Pima Indians Diabetes Database, 2016. URL <https://www.kaggle.com/uciml/pima-indians-diabetes-database>.
- Robert Detrano. V.A. Medical Center, Long Beach and Cleveland Clinic Foundation: Robert Detrano, MD, Ph.D., Donor: David W. Aha, 1998. URL <https://archive.ics.uci.edu/ml/datasets/Heart+Disease>.
- Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.

Supplementary Material: ‘Online Algorithm for Unsupervised Sequential Selection with Contextual Information’

A Missing proofs from Section 2

A.1 Proof of Lemma 1

Lemma 1. For any i, j , and $x_t \in \mathcal{X}$, $\gamma_i(x_t) - \gamma_j(x_t) = p_{ij}^{(t)} - 2\mathbb{P}\{Y_t^i = Y_t, Y_t^j \neq Y_t | X = x_t\}$.

Proof. Using definition of $\gamma_i(x_t) \doteq \mathbb{P}\{Y_t^i \neq Y_t | X = x_t\}$, we get

$$\gamma_i(x_t) - \gamma_j(x_t) = \mathbb{P}\{Y_t^i \neq Y_t | X = x_t\} - \mathbb{P}\{Y_t^j \neq Y_t | X = x_t\}.$$

As the observed feedback is binary, if $Y_t^i = Y_t^j$ and $Y_t^i \neq Y_t$ then $Y_t^j \neq Y_t$,

$$\begin{aligned} \gamma_i(x_t) - \gamma_j(x_t) &= \mathbb{P}\left\{\cancel{Y_t^i \neq Y_t, Y_t^i = Y_t^j} | X = x_t\right\} + \mathbb{P}\{Y_t^i \neq Y_t, Y_t^i \neq Y_t^j | X = x_t\} \\ &\quad - \mathbb{P}\left\{\cancel{Y_t^j \neq Y_t, Y_t^i = Y_t^j} | X = x_t\right\} - \mathbb{P}\{Y_t^j \neq Y_t, Y_t^i \neq Y_t^j | X = x_t\}. \end{aligned}$$

Adding and subtracting $\mathbb{P}\{Y_t^i = Y_t, Y_t^i \neq Y_t^j | X = x_t\}$,

$$\begin{aligned} \gamma_i(x_t) - \gamma_j(x_t) &= \mathbb{P}\{Y_t^i \neq Y_t, Y_t^i \neq Y_t^j | X = x_t\} + \mathbb{P}\{Y_t^i = Y_t, Y_t^i \neq Y_t^j | X = x_t\} \\ &\quad - \mathbb{P}\{Y_t^j \neq Y_t, Y_t^i \neq Y_t^j | X = x_t\} - \mathbb{P}\{Y_t^i = Y_t, Y_t^i \neq Y_t^j | X = x_t\}. \end{aligned}$$

If $Y_t^i \neq Y_t^j$ and $Y_t^j \neq Y_t$ then $Y_t^i = Y_t$,

$$\begin{aligned} \gamma_i(x_t) - \gamma_j(x_t) &= \mathbb{P}\{Y_t^i \neq Y_t^j | X = x_t\} - \mathbb{P}\{Y_t^i = Y_t, Y_t^i \neq Y_t^j | X = x_t\} \\ &\quad - \mathbb{P}\{Y_t^i = Y_t, Y_t^i \neq Y_t^j | X = x_t\} \\ &= \mathbb{P}\{Y_t^i \neq Y_t^j | X = x_t\} - 2\mathbb{P}\{Y_t^i = Y_t, Y_t^i \neq Y_t^j | X = x_t\}. \\ \implies \gamma_i(x_t^i) - \gamma_j(x_t^j) &= p_{ij}^{(t)} - 2\mathbb{P}\{Y_t^i = Y_t, Y_t^j \neq Y_t | X = x_t\}. \quad \square \end{aligned}$$

A.2 Proof of Lemma 2

Lemma 2. Let $P \in \mathcal{P}_{\text{CWD}}$ and $\mathcal{B}_t = \{i : \forall j > i, C_j - C_i > p_{ij}^{(t)}\} \cup \{K\}$. Then the arm $I_t = \min(\mathcal{B}_t)$ is the optimal arm for a context x_t .

Proof. Let i_t^* be an optimal arm for a context x_t . As $p_{ij}^{(t)} = \mathbb{P}\{Y_t^i \neq Y_t^j | X = x_t\}$ and i_t^* is an optimal arm, we have $\forall j < i_t^* : C_{i_t^*} - C_j \leq \mathbb{P}\{Y_t^{i_t^*} \neq Y_t^j | X = x_t\} \implies C_{i_t^*} - C_j \not> \mathbb{P}\{Y_t^{i_t^*} \neq Y_t^j | X = x_t\} \implies \forall j < i_t^* \notin \mathcal{B}_t$. If any sub-optimal arm $h \in \mathcal{B}_t$ then $h > i_t^*$ i.e.,

$$\mathcal{B}_t = \{i_t^*, h_1, \dots, h_n, K\},$$

where $i_t^* < h_1 < \dots < h_n < K$. By construction of set $\mathcal{B}_t$, the minimum indexed arm in set $\mathcal{B}_t$ is only the optimal arm. $\square$

A.3 Proof of Theorem 1

We need the following results to proof of Theorem 1.

Lemma 7. Let $i < j$ and $x_t \in \mathcal{X}$ be any context. Assume

$$C_j - C_i \notin \left(\gamma_i(x_t) - \gamma_j(x_t), \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\} \right]. \quad (11)$$

Then, $C_j - C_i > \gamma_i(x_t) - \gamma_j(x_t)$ iff $C_j - C_i > \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$.

Proof. Assume that $C_j - C_i > \gamma_i(x_t) - \gamma_j(x_t)$. As $C_j - C_i \notin \left(\gamma_i(x_t) - \gamma_j(x_t), \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\} \right]$, we get $C_j - C_i > \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$. The proof of other direction follows by noting that $\mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\} \geq \gamma_i(x_t) - \gamma_j(x_t)$. $\square$

Lemma 8. Let $i > j$ and $x_t \in \mathcal{X}$ be any context. Assume

$$C_i - C_j \notin \left(\gamma_j(x_t) - \gamma_i(x_t), \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\} \right]. \quad (12)$$

Then, $C_i - C_j \leq \gamma_j(x_t) - \gamma_i(x_t)$ iff $C_j - C_i \leq \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$.

Proof. Let $C_i - C_j \leq \gamma_j(x_t) - \gamma_i(x_t)$. As $\gamma_j(x_t) - \gamma_i(x_t) \leq \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$, we get $C_i - C_j \leq \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$.

The condition $C_i - C_j \leq \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$ along with $C_i - C_j \notin \left(\gamma_j(x_t) - \gamma_i(x_t), \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\} \right]$ implies the other direction, i.e., $C_i - C_j \leq \gamma_j(x_t) - \gamma_i(x_t)$. $\square$

Lemma 9. Let i_t^* be an optimal arm for a context x_t . Any problem instance $P \in \mathcal{P}_{USS}$ is learnable if for every context in P following holds:

$$\forall j > i_t^*, C_j - C_{i_t^*} > \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\}.$$

The proof of Lemma 9 follows from Lemma 7 and Lemma 8. Now we give proof for Theorem 1.

Theorem 1. The set $\mathcal{P}_{CWD}$ is maximal learnable.

Proof. Let i_t^* be an optimal arm for a context x_t . It is enough to prove that any problem instance $P \in \mathcal{P}_{USS}$ is learnable if

$$\forall j > i_t^*, C_j - C_{i_t^*} > \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\}. \quad (\text{definition of CWD property})$$

From Lemma 7 and Lemma 8, if the optimal arm satisfies following conditions,

$$\begin{aligned} \forall j > i_t^*, C_j - C_{i_t^*} &\notin \left(\gamma_{i_t^*}(x_t) - \gamma_j(x_t), \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\} \right] \text{ and} \\ \forall j < i_t^*, C_{i_t^*} - C_j &\notin \left(\gamma_j(x_t) - \gamma_{i_t^*}(x_t), \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\} \right], \end{aligned}$$

then, for $j > i_t^*$, $C_j - C_{i_t^*} > \gamma_{i_t^*}(x_t) - \gamma_j(x)$ iff $C_j - C_{i_t^*} > \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\}$ and for $j < i_t^*$, $C_{i_t^*} - C_j \leq \gamma_j(x) - \gamma_{i_t^*}(x_t)$ iff $C_j - C_{i_t^*} \leq \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\}$. Hence we can use $\mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}$ as a proxy for $\gamma_{i_t^*}(x) - \gamma_j(x)$ to make decision about the optimal arm. Now notice that for $j < i_t^*$, $C_{i_t^*} - C_j \leq \gamma_j(x) - \gamma_{i_t^*}(x_t)$. Hence,

$$\forall j < i_t^*, C_{i_t^*} - C_j \notin \left(\gamma_j(x_t) - \gamma_{i_t^*}(x_t), \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\} \right] \text{ and}$$

$$\forall j > i_t^*, C_j - C_{i_t^*} \notin \left(\gamma_{i_t^*}(x_t) - \gamma_j(x_t), \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\} \right) \quad (13)$$

are sufficient for learnability. Note that Eq. (13) is equivalent to

$$\forall j > i_t^*, C_j - C_{i_t^*} > \mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\}. \quad (14)$$

Note that if Eq. (14) does not hold, then knowing $\mathbb{P} \left\{ Y_t^{i_t^*} \neq Y_t^j | X = x_t \right\}$ is not sufficient for finding the optimal arm. $\square$

A.4 Regret decomposition when contexts satisfy WD with some known probability

Without knowing the disagreement probability, it is impossible to check whether a context satisfies WD property or not. Hence we consider a case where a context can satisfy WD property with some fixed probability. For such cases, we can decompose the regret into two parts: regret due to the contexts that satisfy WD property and regret due to the contexts that do not satisfy WD property. Note that the regret can be linear due to the contexts that do not satisfy the WD condition.

Our next result gives the upper bound on the regret where the contexts satisfy WD property with a known fixed probability.

Lemma 10. *Let ρ be the probability of context that it does not satisfy the WD property and R_{max} be the maximum regret incurred for any context. If $\mathfrak{R}_T$ is the regret incurred when all contexts satisfy WD property then, the regret incurred when contexts satisfy WD with probability $(1 - \rho)$ is given by*

$$\mathfrak{R}'_T \leq (1 - \rho)\mathfrak{R}_T + \rho R_{max}T.$$

Proof. Let ρ be the probability of context that it does not satisfy the WD property and $r_t(I_t, i_t^*)$ be the regret incurred for selecting sub-optimal arm I_t for the context x_t . Then the regret can be decomposed into two parts as follows:

$$\begin{aligned} \mathfrak{R}'_T &= \mathbb{E} \left[\sum_{t=1}^T \left[\mathbb{1}_{\{x_t \text{ satisfies WD}\}} r_t(I_t, i_t^*) + \mathbb{1}_{\{x_t \text{ does not satisfy WD}\}} r_t(I_t, i_t^*) \right] \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}_{\{x_t \text{ satisfies WD}\}} r_t(I_t, i_t^*) \right] + \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}_{\{x_t \text{ does not satisfy WD}\}} r_t(I_t, i_t^*) \right] \\ &= \sum_{t=1}^T \mathbb{P} \{x_t \text{ satisfies WD}\} r_t(I_t, i_t^*) + \sum_{t=1}^T \mathbb{P} \{x_t \text{ does not satisfy WD}\} r_t(I_t, i_t^*). \end{aligned} \quad (15)$$

First, we will bound the regret due to the contexts that do not satisfy WD property (second term of Eq. (15)). Note that the context that does not satisfy WD property, the learner can not make the correct decision hence always incurs regret. Since the maximum regret is upper bounded by R_{max} , we have

$$\sum_{t=1}^T \mathbb{P} \{x_t \text{ does not satisfy WD}\} r_t(I_t, i_t^*) \leq \sum_{t=1}^T \mathbb{P} \{x_t \text{ does not satisfy WD}\} R_{max}$$

Since ρ is the probability of context that it does not satisfy the WD property, we get

$$\begin{aligned} &= \sum_{t=1}^T \rho R_{max} \\ \implies \sum_{t=1}^T \mathbb{P} \{x_t \text{ does not satisfy WD}\} r_t(I_t, i_t^*) &\leq \rho R_{max}T. \end{aligned} \quad (16)$$

Now we will bound the regret due to the contexts which satisfy WD property (first term in Eq. (15)). Since any context satisfies WD with $1 - \rho$ probability, we have

$$\sum_{t=1}^T \mathbb{P} \{x_t \text{ satisfies WD}\} r_t(I_t, i_t^*) = \sum_{t=1}^T (1 - \rho) r_t(I_t, i_t^*) = (1 - \rho) \sum_{t=1}^T r_t(I_t, i_t^*). \quad (17)$$

By assuming that all contexts are satisfying WD property, we have regret $\mathfrak{R}_T = \sum_{t=1}^T r_t(I_t, i_t^*)$. Using it with Eq. (16) in Eq. (17), we get

$$\mathfrak{R}'_T \leq (1 - \rho)\mathfrak{R}_T + \rho R_{\max} T. \quad \square$$

B Missing proofs from Section 4

B.1 Proof of Lemma 3

Lemma 3. Let $V_{ij}^t = \sum_{s \in S_{ij}^t} \Phi_{ij}(x_s) \Phi_{ij}(x_s)^\top$, $\Sigma_{ij} = \mathbb{E} [\Phi_{ij}(X) \Phi_{ij}(X)^\top]$, Ψ and $\delta \in (0, 1)$ be two positive constants. Then, there exist positive universal constants C_1 and C_2 such that the minimum eigenvalue of $\lambda_{\min}(V_{ij}^t) \geq \Psi$ with probability at least $1 - 2\delta/K^2$, iff

$$|S_{ij}^t| \geq \left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log(K^2/2\delta)}}{\lambda_{\min}(\Sigma_{ij})} \right)^2 + \frac{2\Psi}{\lambda_{\min}(\Sigma_{ij})}.$$

Proof. The result is adapted from (Li et al., 2017, Proposition 1), which uses the standard random matrix theory result from (Vershynin, 2012, Theorem 5.39). We need to carefully construct the sample complexity bound for our case as the observations are only observed for a pair of arms. $\square$

The following result is needed to prove Lemma 4.

Lemma 11. Let $\bar{V}_{ij}^t = \lambda I_{d'} + a \sum_{s \in S_{ij}^t} \Phi_{ij}(x_s) \Phi_{ij}(x_s)^\top$ for any (i, j) pair of arms, and $n_{ij}^t = |S_{ij}^t|$. Then

$$\det(\bar{V}_{ij}^t) \leq (\lambda + a n_{ij}^t / d')^{d'}.$$

Proof. The proof is adapted from Lemma 10 of Abbasi-Yadkori et al. (2011). By using inequality of arithmetic and geometric means, we have $\det(\bar{V}_{ij}^t) \leq (\text{trace}(\bar{V}_{ij}^t)/d')^{d'}$. As the trace of matrix is a linear mapping i.e. $\text{trace}(A + B) = \text{trace}(A) + \text{trace}(B)$, hence, we get

$$\begin{aligned} \text{trace}(\bar{V}_{ij}^t) &= \text{trace}(\lambda I_{d'}) + a \sum_{s \in S_{ij}^t} \text{trace}(\Phi_{ij}(x_s) \Phi_{ij}(x_s)^\top) \\ &= \lambda d' + a \sum_{s \in S_{ij}^t} \|\Phi_{ij}(x_s)\|_2^2 \\ &\leq \lambda d' + a n_{ij}^t. \quad (\text{as } \|\Phi_{ij}(x)\|_2 \leq 1 \text{ and } n_{ij}^t = |S_{ij}^t|) \end{aligned}$$

Using upper bound of $\text{trace}(\bar{V}_{ij}^t)$ for bounding $\det(\bar{V}_{ij}^t)$, we get

$$\det(\bar{V}_{ij}^t) \leq (\text{trace}(\bar{V}_{ij}^t)/d')^{d'} \leq ((\lambda d' + a n_{ij}^t)/d')^{d'} \leq (\lambda + a n_{ij}^t / d')^{d'}. \quad \square$$

B.2 Proof of Lemma 4

Lemma 4 (Confidence Ellipsoid). Let m be such that $\lambda_{\min}(V_{ij}^{m+1}) \geq 1$ for any pair (i, j) . Then the following event holds with probability at least $1 - 2\delta/K^2$ for USS-PD:

$$\|\hat{\theta}_{ij}^t - \theta_{ij}^*\|_{V_{ij}^t} \leq \alpha_{ij}^t, \quad \forall t > m$$

where $\alpha_{ij}^t = \frac{2\sigma}{\kappa} \sqrt{\frac{d'}{2} \log(1 + \frac{2t}{d'}) + \log(\frac{K^2}{2\delta})}$.

Proof. Let $\bar{V}_{ij}^t = \lambda I_{d'} + V_{ij}^t$. If Eq. (9) is used for estimation of unknown parameter θ_{ij}^* , then by using Eq. (26) and Lemma 8 of Li et al. (2017) with $\lambda_{\min}(V_{ij}^{m+1}) \geq 1$, we have

$$\|\hat{\theta}_{ij}^t - \theta_{ij}^*\|_{V_{ij}^t} \leq \frac{1}{\kappa} \left\| \sum_{s \in S_{ij}^t} \varepsilon_s \Phi_{ij}(x_s) \right\|_{(V_{ij}^t)^{-1}} \leq \frac{(1 - \lambda)^{-\frac{1}{2}}}{\kappa} \left\| \sum_{s \in S_{ij}^t} \varepsilon_s \Phi_{ij}(x_s) \right\|_{(\bar{V}_{ij}^t)^{-1}}.$$

Using upper bound of $\left\| \sum_{s \in S_{ij}^t} \varepsilon_s \Phi_{ij}(x_s) \right\|_{(\bar{V}_{ij}^t)^{-1}}$ as given in Theorem 1 of [Abbasi-Yadkori et al. \(2011\)](#) where ε_s is σ -subGaussian random variable, the following inequality holds with at least probability $1 - 2\delta/K^2$

$$\begin{aligned} &\leq \frac{(1-\lambda)^{-\frac{1}{2}}}{\kappa} \sqrt{2\sigma^2 \log \left(\frac{\det(\bar{V}_{ij}^t)^{1/2} \det(\lambda \mathbf{I}_{d'})^{-1/2}}{2\delta/K^2} \right)} \\ &= \frac{\sigma(1-\lambda)^{-\frac{1}{2}}}{\kappa} \sqrt{2 \log \left(\frac{\det(\bar{V}_{ij}^t)}{\det(\lambda \mathbf{I}_{d'})} \right)^{\frac{1}{2}} + 2 \log \left(\frac{K^2}{2\delta} \right)}. \end{aligned}$$

Upper bounding $\det(\bar{V}_{ij}^t)$ with $a = 1$, $\lambda = 1/2$, and $n_{ij}^t \leq t$ by using Lemma 11, we get

$$\begin{aligned} \Rightarrow \left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{V_{ij}^t} &\leq \frac{2\sigma}{\kappa} \sqrt{\frac{d'}{2} \log \left(1 + \frac{2n_{ij}^t}{d'} \right) + \log \left(\frac{K^2}{2\delta} \right)} \\ &\leq \frac{2\sigma}{\kappa} \sqrt{\frac{d'}{2} \log \left(1 + \frac{2t}{d'} \right) + \log \left(\frac{K^2}{2\delta} \right)}. \quad \square \end{aligned}$$

B.3 Proof of Lemma 5

Lemma 5. Let $\theta \in \Theta_{\text{CWD}}$. Then **USS-PD** prefers any sub-optimal arm $l < i_t^*$ for context x_t with probability at most $\delta/2$.

Proof. If sub-optimal arm $l < i_t^*$ is preferred by **USS-PD** then using Eq. (10b), we get

$$\begin{aligned} \mathbb{1}_{\{l \succ_t i_t^*, i_t^* = i\}} &= \mathbb{1}_{\{C_i - C_l > \tilde{p}_{li}^{(t)}, I_t = l, i_t^* = i\}} \\ &\leq \mathbb{1}_{\{C_i - C_l > \tilde{p}_{li}^{(t)}\}}. \quad (\text{as } A \cap B \cap C \subseteq A) \end{aligned}$$

Using $C_i - C_l = p_{li}(x_t) - \xi_{li}(x_t)$ for $l < i$, we have

$$\Rightarrow \mathbb{1}_{\{l \succ_t i_t^*, i_t^* = i\}} = \mathbb{1}_{\{p_{li}(x_t) - \xi_{li}(x_t) > \tilde{p}_{li}^{(t)}\}} = \mathbb{1}_{\{p_{li}(x_t) - \tilde{p}_{li}^{(t)} > \xi_{li}(x_t)\}}.$$

Using definition of $p_{li}(x_t)$ and $\tilde{p}_{li}^{(t)}$,

$$\Rightarrow \mathbb{1}_{\{l \succ_t i_t^*, i_t^* = i\}} = \mathbb{1}_{\left\{ \mu(\Phi_{li}(x_t)^\top \theta_{li}^*) - \mu(\Phi_{li}(x_t)^\top \hat{\theta}_{li}^t + \alpha_{li}^t \|\Phi_{li}(x_t)\|_{(V_{li}^t)^{-1}}) > \xi_{li}(x_t) \right\}}.$$

Since $\mu(\cdot)$ is an increasing function and using α_{li}^t as defined in Lemma 4, $\mu(\Phi_{li}(x_t)^\top \hat{\theta}_{li}^t + \alpha_{li}^t \|\Phi_{li}(x_t)\|_{(V_{li}^t)^{-1}})$ is the upper bound on $\mu(\Phi_{li}(x_t)^\top \theta_{li}^*)$ for all (l, i) pairs with probability at least $1 - \delta/2$. We show it as follows:

$$\begin{aligned} \Phi_{li}(x_t)^\top \theta_{li}^* &= \Phi_{li}(x_t)^\top \hat{\theta}_{li}^t + \Phi_{li}(x_t)^\top (\theta_{li}^* - \hat{\theta}_{li}^t) \\ &= \Phi_{li}(x_t)^\top \hat{\theta}_{li}^t + \|\Phi_{li}(x_t)\|_{(V_{li}^t)^{-1}} \left\| \theta_{li}^* - \hat{\theta}_{li}^t \right\|_{V_{li}^t} \\ \Rightarrow \Phi_{li}(x_t)^\top \theta_{li}^* &\leq \Phi_{li}(x_t)^\top \hat{\theta}_{li}^t + \alpha_{li}^t \|\Phi_{li}(x_t)\|_{(V_{li}^t)^{-1}}. \quad \left(\text{using } \left\| \theta_{li}^* - \hat{\theta}_{li}^t \right\|_{V_{li}^t} \leq \alpha_{li}^t \right) \end{aligned}$$

Since $\mu(\cdot)$ is an increasing function,

$$\Rightarrow \mu(\Phi_{li}(x_t)^\top \theta_{li}^*) \leq \mu(\Phi_{li}(x_t)^\top \hat{\theta}_{li}^t + \alpha_{li}^t \|\Phi_{li}(x_t)\|_{(V_{li}^t)^{-1}}).$$

Hence, any sub-optimal arm smaller than the optimal arm is selected by **USS-PD** with probability at most $\delta/2$. It completes the proof of the lemma. $\square$

B.4 Proof of Lemma 6

Lemma 6. Let $\theta \in \Theta_{\text{CWD}}$. If **USS-PD** prefers a sub-optimal arm $h > i_t^*$ for context x_t then

$$2k_\mu \alpha_{i_t^* h}^t > \xi_{i_t^* h}(x_t) \sqrt{\lambda_{\min}(V_{i_t^* h}^t)}.$$

where $\xi_{i_t^* h} = C_h - C_{i_t^*} - p_{i_t^* h}^{(t)}$ and $\alpha_{i_t^* h}^t$ is given by Lemma 4.

Proof. If sub-optimal arm $h > i_t^*$ is preferred by **USS-PD** then using Eq. (10a), we get

$$\begin{aligned} \mathbb{1}_{\{h \succ_t i, i_t^* = i\}} &= \mathbb{1}_{\{C_h - C_i < \bar{p}_{ih}^{(t)}, h \succ_t i_t^*, i_t^* = i\}} \\ &\leq \mathbb{1}_{\{C_h - C_i < \bar{p}_{ih}^{(t)}\}}. \quad (\text{as } A \cap B \cap C \subseteq A) \end{aligned}$$

Using $C_h - C_i = p_{ih}(x_t) + \xi_{ih}(x_t)$ for $h > i$, we get

$$\implies \mathbb{1}_{\{h \succ_t i, i_t^* = i\}} = \mathbb{1}_{\{p_{ih}(x_t) + \xi_{ih}(x_t) < \bar{p}_{ih}^{(t)}\}} = \mathbb{1}_{\{\bar{p}_{ih}^{(t)} - p_{ih}(x_t) > \xi_{ih}(x_t)\}}.$$

Using definition of $p_{ih}(x_t)$ and $\bar{p}_{ih}^{(t)}$,

$$\implies \mathbb{1}_{\{h \succ_t i, i_t^* = i\}} = \mathbb{1}_{\left\{ \mu \left(\Phi_{ih}(x_t)^\top \theta_{ih}^* + \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} \right) - \mu \left(\Phi_{ih}(x_t)^\top \hat{\theta}_{ih}^t \right) > \xi_{ih}(x_t) \right\}}.$$

As μ is Lipschitz, $|\mu(z_1) - \mu(z_2)| \leq k_\mu |z_1 - z_2|$ where k_μ is Lipschitz constant, we have

$$\begin{aligned} &\leq \mathbb{1}_{\left\{ k_\mu \|\Phi_{ih}(x_t)^\top \theta_{ih}^* + \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} - \Phi_{ih}(x_t)^\top \hat{\theta}_{ih}^t \| > \xi_{ih}(x_t) \right\}} \\ &\leq \mathbb{1}_{\left\{ k_\mu \|\Phi_{ih}(x_t)^\top \theta_{ih}^* - \Phi_{ih}(x_t)^\top \hat{\theta}_{ih}^t\| + k_\mu \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} > \xi_{ih}(x_t) \right\}} \\ &= \mathbb{1}_{\left\{ k_\mu \|\Phi_{ih}(x_t)^\top (\theta_{ih}^* - \hat{\theta}_{ih}^t)\| + k_\mu \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} > \xi_{ih}(x_t) \right\}}. \end{aligned}$$

Using Cauchy-Schwartz inequality and $\|x\|_A^2 = x^\top A x$, we get

$$\leq \mathbb{1}_{\left\{ k_\mu \|\Phi_{ih}\|_{(V_{ih}^t)^{-1}} \|\theta_{ih}^* - \hat{\theta}_{ih}^t\|_{V_{ih}^t} + k_\mu \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} > \xi_{ih}(x_t) \right\}}.$$

As $\|\theta_{ih}^* - \hat{\theta}_{ih}^t\|_{V_{ih}^t} \leq \alpha_{ih}^t$, we get

$$\begin{aligned} &\leq \mathbb{1}_{\left\{ k_\mu \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} + k_\mu \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} > \xi_{ih}(x_t) \right\}} \\ &= \mathbb{1}_{\left\{ 2k_\mu \alpha_{ih}^t \|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} > \xi_{ih}(x_t) \right\}}. \end{aligned}$$

As $\|\Phi_{ih}(x_t)\|_{(V_{ih}^t)^{-1}} \leq \|\Phi_{ih}(x_t)\|_2 / \sqrt{\lambda_{\min}(V_{ih}^t)}$ where $\lambda_{\min}(V_{ih}^t)$ is the smallest eigenvalue of matrix V_{ih}^t and $\|\Phi_{ih}(x_t)\|_2 \leq 1$, we get

$$\implies \mathbb{1}_{\{h \succ_t i, i_t^* = i\}} \leq \mathbb{1}_{\left\{ 2k_\mu \alpha_{ih}^t > \xi_{ih}(x_t) \sqrt{\lambda_{\min}(V_{ih}^t)} \right\}}. \quad (18)$$

The event on LHS is subset of event of RHS in Eq. (18). By changing i to i_t^* completes the proof of the lemma. $\square$

B.5 Proof of Theorem 2

Theorem 2 (Regret Upper Bound). Let $\theta \in \Theta_{\text{CWD}}$, $\delta \in (0, 1)$, Assumption 1 holds, and $\xi_h = \min_{t \geq 1} \xi_{i_t^* h}(x_t)$. Then with probability at least $1 - 2\delta$, the regret of **USS-PD** for $T > m$ contexts is

$$\begin{aligned} \mathfrak{R}_T &\leq R_{\max} \left[m + \sum_{h=2}^K \left(\left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log \left(\frac{K^2}{2\delta} \right)}}{\lambda_\Sigma} \right)^2 + \frac{16}{\lambda_\Sigma} \right. \right. \\ &\quad \left. \left. \left(\frac{k_\mu \sigma}{\xi_h \kappa} \right)^2 \left(\frac{d'}{2} \log \left(1 + \frac{2T}{d'} \right) + \log \left(\frac{K^2}{2\delta} \right) \right) \right) \right]. \end{aligned}$$

Proof. The regret for T rounds in the Contextual USS problem is given by

$$\mathfrak{R}_T = \sum_{t=1}^T (C_{I_t} + \gamma_{I_t}(x_t) - (C_{i_t^*} + \gamma_{i_t^*}(x_t))).$$

As R_{max} denote the maximum regret incurred for any context, we get

$$\mathfrak{R}_T \leq R_{max} \sum_{t=1}^T \mathbb{1}_{\{I_t \neq i_t^*\}}. \quad (19)$$

As $\mathbb{1}_{\{I_t \neq i_t^*\}}$ has two random quantities I_t and i_t^* , we can re-write it as follows:

$$\mathbb{1}_{\{I_t \neq i_t^*\}} = \sum_{l < i} \mathbb{1}_{\{I_t=l, i_t^*=i\}} + \sum_{h' > i} \mathbb{1}_{\{I_t=h', i_t^*=i\}}.$$

Note that if **USS-PD** selects $l < i_t^*$ then l must be preferred over i_t^* whereas if $h' > i_t^*$ is selected then there exists an arm $h > i_t^*$ which is preferred over i_t^* . Hence, we have

$$\begin{aligned} \mathbb{1}_{\{I_t \neq i_t^*\}} &= \sum_{l < i} \mathbb{1}_{\{l \succ_t i_t^*, i_t^*=i\}} + \sum_{h' > i} \mathbb{1}_{\{I_t=h', h \succ_t i_t^*, i_t^*=i\}} \\ &\leq \sum_{l < i} \mathbb{1}_{\{l \succ_t i_t^*, i_t^*=i\}} + \sum_{h > i} \mathbb{1}_{\{h \succ_t i_t^*, i_t^*=i\}}. \end{aligned} \quad (20)$$

Using above bound in Eq. (19), we get

$$\mathfrak{R}_T \leq R_{max} \sum_{t=1}^T \left[\sum_{l < i} \mathbb{1}_{\{l \succ_t i_t^*, i_t^*=i\}} + \sum_{h > i} \mathbb{1}_{\{h \succ_t i_t^*, i_t^*=i\}} \right].$$

From Lemma 5, $\mathbb{1}_{\{l \succ_t i_t^*, i_t^*=i\}} = 0$ for any $l < i$ with probability at least $1 - \delta/2$, then the regret becomes

$$\mathfrak{R}_T \leq R_{max} \sum_{t=1}^T \sum_{h > i} \mathbb{1}_{\{h \succ_t i_t^*, i_t^*=i\}} = R_{max} \sum_{h > i} \sum_{t=1}^T \mathbb{1}_{\{h \succ_t i_t^*, i_t^*=i\}} \leq R_{max} \sum_{h=2}^K \sum_{t=1}^T \mathbb{1}_{\{h \succ_t i_t^*, i_t^* < h\}}.$$

Note that α_{ih}^t is slowly increasing value with t that implies $\alpha_{ih}^t \leq \alpha_{ih}^T$ for all $t \leq T$. Using Lemma 3 with $\Psi = \left(\frac{2k_\mu \alpha_{ih}^T}{\xi_{ih}} \right)^2$, $\Sigma_{ih} = \mathbb{E} [\Phi_{ih}(X_s) \Phi_{ih}(X_s)^T]$ where $s \in S_{ih}^t$, after

$$n_{ih}^T \doteq \left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log(K^2/2\delta)}}{\lambda_{min}(\Sigma_{ih})} \right)^2 + \frac{2}{\lambda_{min}(\Sigma_{ih})} \left(\frac{2k_\mu \alpha_{ih}^T}{\xi_{ih}} \right)^2$$

observations for arm pair (i, h) the $\lambda_{min}(V_{ih}^t) \geq \left(\frac{2k_\mu \alpha_{ih}^T}{\xi_{ih}} \right)^2$ with probability at least $1 - 2\delta/K^2$.

Therefore, after having n_{ih}^T observations, the sub-optimal arm $h(> i)$ will not be preferred over optimal arm i with probability at least $1 - 2\delta/K^2$. Therefore, with probability at least $1 - 2\delta/K^2$, following equations also hold

$$\mathbb{1}_{\{I_t=h, i_t^*=i, |S_{ih}^t| \geq n_{ih}^T\}} = 0 \implies \sum_{t=1}^T \mathbb{1}_{\{I_t=h, i_t^*=i, h > i\}} \leq n_{ih}^T.$$

Due to the problem structure, whenever an arm h is selected, disagreement labels for all arm pair (i, j) where $i < j \leq h$ are observed. Therefore, with probability at least $1 - 2\delta/K$ (by union bound), the maximum number of times an arm h is selected when the optimal arm's index is smaller than h is n_h^T such that

$$n_h^T = \left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log(K^2/2\delta)}}{\lambda_\Sigma} \right)^2 + \frac{2}{\lambda_\Sigma} \left(\frac{2k_\mu \alpha_T}{\xi_h} \right)^2$$

$$= \left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log(K^2/2\delta)}}{\lambda_\Sigma} \right)^2 + \frac{8}{\lambda_\Sigma} \left(\frac{k_\mu \alpha_T}{\xi_h} \right)^2$$

where $\xi_h = \min_{i < h, t \geq 1} \xi_{ih}(x_t)$, $\lambda_\Sigma = \min_{i < j \leq K} \lambda_{\min}(\mathbb{E}[\Phi_{ij}(X_s)\Phi_{ij}(X_s)^\top])$ and $\alpha_T \geq \max_{i < h} \alpha_{ih}^T$. By using union bound, we get following bound with probability at least $1 - \delta/2K$

$$\sum_{t=1}^T \mathbb{1}_{\{I_t=h, i_t^* < h\}} \leq n_h^T. \quad (21)$$

From Eq. (21), using $\sum_{t=1}^T \mathbb{1}_{\{I_t=h, i_t^* < h\}} \leq n_h^T$ and value of n_h^T , we get following upper bound on regret that holds with probability at least $1 - \delta$ by union bound

$$\mathfrak{R}_T \leq R_{\max} \sum_{h=2}^K n_h^T = R_{\max} \sum_{h=2}^K \left(\left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log(K^2/2\delta)}}{\lambda_\Sigma} \right)^2 + \frac{8}{\lambda_\Sigma} \left(\frac{k_\mu \alpha_T}{\xi_h} \right)^2 \right).$$

Using $\alpha_T = \frac{2\sigma}{\kappa} \sqrt{\frac{d'}{2} \log(1 + 2T/d') + \log(K^2/2\delta)}$ from Lemma 4 that ensures parameter θ_{ij}^* bounds for all pairs (i, j) holds with probability at least $1 - \delta/2K$ (by union bound) for $T > m$ where $m = C\lambda_\Sigma^{-2} (d + \log(k^2/2\delta)) + 2\lambda_\Sigma^{-1}$ such that $\lambda_{\min}(V_{ij}^{m+1}) \geq 1$ for all pair (i, j) , we have

$$\begin{aligned} \mathfrak{R}_T &\leq R_{\max} \left(m + \sum_{h=2}^K n_h^T \right) \\ \implies \mathfrak{R}_T &\leq R_{\max} \left[m + \sum_{h=2}^K \left(\left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log(K^2/2\delta)}}{\lambda_\Sigma} \right)^2 + \frac{16}{\lambda_\Sigma} \left(\frac{k_\mu \sigma}{\xi_h \kappa} \right)^2 \left(\frac{d'}{2} \log\left(1 + \frac{2T}{d'}\right) + \log\left(\frac{K^2}{2\delta}\right) \right) \right) \right]. \quad \square \end{aligned}$$

B.6 Algorithm with Regularization Term

USS-PD uses forced exploration by selecting arm K until the correlation matrix V_{ij}^t is not invertible for all (i, j) pairs of arms. Further, the minimum eigenvalue of V_{ij}^t for all (i, j) pairs is needed to be larger than 1 so that bound given in Lemma 4 holds. Alternatively, V_{ij}^t can be initialized by adding a regularization term (Abbasi-Yadkori et al., 2011; Zhang et al., 2016; Jun et al., 2017) to avoid forced exploration and then apply OFUL type analysis. We have given an algorithm named **USS-PD- λ I** which uses regularization term $\lambda I_{d'}$. However, its analysis still needed the minimum eigenvalue of the non-regularized part of the correlation matrix to become larger than some positive value (depends on λ value), as shown in our next result.

USS-PD- λ I Algorithm for Contextual USS using Pairwise Disagreement with λ I Initialization

- 1: **Input:** Tuning parameters: $\delta \in (0, 1)$ and $\lambda > 0$
 - 2: Select arm K for first context x_1
 - 3: $\forall i < j \leq K$: set $\bar{V}_{ij}^1 \leftarrow \lambda I_{d'} + \Phi_{ij}(x_1)\Phi_{ij}(x_1)^\top$ and update $\hat{\theta}_{ij}^1$ by solving Eq. (9)
 - 4: **for** $t = 2, 3, \dots$ **do**
 - 5: Receive context x_t . Set $i = 1$ and $I_t = 0$
 - 6: **do**
 - 7: Play arm i
 - 8: $\forall j \in [i + 1, K]$: compute $\bar{p}_{ij}^{(t)} \leftarrow \mu \left(\Phi_{ij}(x_t)^\top \hat{\theta}_{ij}^{t-1} + \alpha_{ij}^{t-1} \|\Phi_{ij}(x_t)\|_{(\bar{V}_{ij}^{t-1})^{-1}} \right)$
 - 9: If $\forall j \in [i + 1, K]$: $C_j - C_i > \bar{p}_{ij}^{(t)}$ or $i = K$ then set $I_t = i$ else set $i = i + 1$
 - 10: **while** $I_t = 0$
 - 11: Select arm I_t and observe $Y_t^1, Y_t^2, \dots, Y_t^{I_t}$
 - 12: $\forall i < j \leq I_t$: update $\bar{V}_{ij}^t \leftarrow \bar{V}_{ij}^{t-1} + \Phi_{ij}(x_t)\Phi_{ij}(x_t)^\top$ and $\hat{\theta}_{ij}^t$ by solving Eq. (9)
 - 13: **end for**
-

Lemma 12. Let $\bar{V}_{ij}^t = \lambda \mathbf{I}_{d'} + V_{ij}^t$ for any $\lambda > 0$ and $\|\theta_{ij}\|_2 \leq S$ for all (i, j) pair. Then for any $t > \min\{s : \forall i < j \ni \lambda_{\min}(V_{ij}(s)) \geq 2\lambda\}$, the following event holds for **USS-PD- λ I** with probability at least $1 - 2\delta/K^2$,

$$\left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{\bar{V}_{ij}^t} \leq \beta_{ij}^t,$$

where $\beta_{ij}^t = \frac{2\sigma}{\kappa} \sqrt{\frac{d'}{2} \log\left(1 + \frac{n_{ij}^t}{d'\lambda}\right) + \log\left(\frac{K^2}{2\delta}\right)} + 2\lambda^{1/2}S$.

Proof. By using $\|Z\|_{A+B} \leq \|Z\|_A + \|Z\|_B$, we have

$$\left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{\bar{V}_{ij}^t} \leq \left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{V_{ij}^t} + \left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{\lambda \mathbf{I}_{d'}}. \quad (\text{as } \bar{V}_{ij}^t = \lambda \mathbf{I}_{d'} + V_{ij}^t)$$

When Eq. (9) is used for estimation of unknown parameter θ_{ij}^* , then by using Eq. (26) and Eq. (27) of Lemma 8 of [Li et al. \(2017\)](#), we have

$$\leq \frac{1}{\kappa} \left\| \sum_{s \in S_{ij}^t} \varepsilon_s \Phi_{ij}(x_s) \right\|_{(V_{ij}^t)^{-1}} + 2\lambda^{1/2}S. \quad (\text{as } \|\theta_{ij}\|_2 \leq S)$$

Sherman Morrison formula gives $\|Z\|_{(V_{ij}^t)^{-1}} \leq \left(1 - \frac{\lambda}{\lambda_{\min}(V_{ij}^t)}\right)^{-\frac{1}{2}} \|Z\|_{(\bar{V}_{ij}^t)^{-1}}$. Using it, we have

$$\left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{\bar{V}_{ij}^t} \leq \frac{\left(1 - \frac{\lambda}{\lambda_{\min}(V_{ij}^t)}\right)^{-\frac{1}{2}}}{\kappa} \left\| \sum_{s \in S_{ij}^t} \varepsilon_s \Phi_{ij}(x_s) \right\|_{(\bar{V}_{ij}^t)^{-1}} + 2\lambda^{1/2}S.$$

Using upper bound on $\left\| \sum_{s \in S_{ij}^t} \varepsilon_s \Phi_{ij}(x_s) \right\|_{(\bar{V}_{ij}^t)^{-1}}$ as given in Theorem 1 of [Abbasi-Yadkori et al. \(2011\)](#), where ε_s is σ -subGaussian random variable and holds with probability at least $1 - 2\delta/K^2$, we get

$$\begin{aligned} \left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{\bar{V}_{ij}^t} &\leq \frac{\left(1 - \frac{\lambda}{\lambda_{\min}(V_{ij}^t)}\right)^{-\frac{1}{2}}}{\kappa} \sqrt{2\sigma^2 \log\left(\frac{\det(\bar{V}_{ij}^t)^{1/2} \det(\lambda \mathbf{I}_{d'})^{-1/2}}{2\delta/K^2}\right)} + 2\lambda^{1/2}S \\ &= \frac{\sigma \left(1 - \frac{\lambda}{\lambda_{\min}(V_{ij}^t)}\right)^{-\frac{1}{2}}}{\kappa} \sqrt{2 \log\left(\frac{\det(\bar{V}_{ij}^t)}{\det(\lambda \mathbf{I}_{d'})}\right)^{\frac{1}{2}} + 2 \log\left(\frac{K^2}{2\delta}\right)} + 2\lambda^{1/2}S. \end{aligned}$$

By using Lemma 11 to upper bound $\det(\bar{V}_{ij}^t)$, where $t > s$ with $a = 1$, and $n_{ij}^t \leq t$, we get

$$\left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{\bar{V}_{ij}^t} \leq \frac{\sigma \left(1 - \frac{\lambda}{\lambda_{\min}(V_{ij}^t)}\right)^{-\frac{1}{2}}}{\kappa} \sqrt{d' \log\left(1 + \frac{n_{ij}^t}{d'\lambda}\right) + 2 \log\left(\frac{K^2}{2\delta}\right)} + 2\lambda^{1/2}S. \quad (22)$$

As $t > s$ such that $\lambda_{\min}(V_{ij}(s)) \geq 2\lambda$, we have

$$\left\| \hat{\theta}_{ij}^t - \theta_{ij}^* \right\|_{\bar{V}_{ij}^t} \leq \frac{2\sigma}{\kappa} \sqrt{\frac{d'}{2} \log\left(1 + \frac{n_{ij}^t}{d'\lambda}\right) + \log\left(\frac{K^2}{2\delta}\right)} + 2\lambda^{1/2}S = \beta_{ij}^t. \quad \square$$

Note that if $\lambda_{\min}(V_{ij}(s)) < \lambda$ then $\left(1 - \frac{\lambda}{\lambda_{\min}(V_{ij}^t)}\right)^{-\frac{1}{2}}$ is not well defined and the bound given in Lemma 12 does not hold. Therefore, $\lambda_{\min}(V_{ij}(s))$ need to be at least greater than λ . Let $m' \doteq C\lambda_{\Sigma}^{-2} (d + \log(k^2/2\delta)) + 4\lambda_{\Sigma}^{-1}\lambda$ where $C > 0$ is the universal constant. Recall $R_{max} \doteq \max_{i \in [K], x \in \mathcal{X}} [C_i + \gamma_i(x) - (C_{i^*} + \gamma_{i^*}(x))]$, where i^* is the optimal arm for a context x . Now we state the regret bounds for **USS-PD- λ I**.

Theorem 3. Let $\theta \in \Theta_{\text{CWD}}$, $\lambda > 0$, $\delta \in (0, 1)$, Assumption 1 holds, and $\xi_h = \min_{t \geq 1} \xi_{i_t^* h}(x_t)$. Then with probability at least $1 - 2\delta$, the regret of **USS-PD- λ** for $T > m'$ contexts is upper bounded as

$$\mathfrak{R}_T \leq R_{\max} \left(m' + \sum_{h=2}^K \left(\left(\frac{C_1 \sqrt{d'} + C_2 \sqrt{\log \left(\frac{K^2}{2\delta} \right)}}{\lambda_\Sigma} \right)^2 + \frac{32\lambda}{\lambda_\Sigma} \left(\frac{k_\mu \sigma}{\xi_h \kappa} \right)^2 \right. \right. \\ \left. \left. \left(\sqrt{\frac{d'}{2} \log \left(1 + \frac{T}{d' \lambda} \right)} + \log \left(\frac{K^2}{2\delta} \right) + 2\lambda^{1/2} S \right)^2 \right) \right).$$

Proof. The proof follows similar steps as Theorem 2 by replacing m by m' and α_{ij}^T by β_{ij}^T . Using $\beta_{ij}^T = \sqrt{\frac{d'}{2} \log \left(1 + \frac{T}{d' \lambda} \right)} + \log \left(\frac{K^2}{2\delta} \right) + 2\lambda^{1/2} S$ completes the proof. $\square$

C Leftover details from Section 5

Since the parameter of each arm (classifier) is known to us (but not to the algorithm), the optimal arm i_t^* can be computed for every context. Therefore, we can also calculate the fraction of contexts for which WD property holds to a given cost vector. To verify WD property for a given context x_t , we first compute disagreement probability for each (i, j) pair of classifiers as³

$$p_{ij}^{(t)} = \mu(x_t^\top \theta_i)(1 - \mu(x_t^\top \theta_j)) + \mu(x_t^\top \theta_j)(1 - \mu(x_t^\top \theta_i)).$$

When all $p_{ij}^{(t)}$ values and i_t^* are known, we can check whether a context x_t satisfies WD property or not by using Eq. (3). For all problem instances derived from the synthetic dataset, the cost vector and the fraction of contexts for which WD property holds are given in Table 1.

PI/Classifiers	Clf. 1	Clf. 2	Clf. 3	Clf. 4	Clf. 5	WD fraction
Costs for PI 1	0.01	0.02	0.032	0.05	0.55	0.997
Costs for PI 2	0.01	0.02	0.032	0.05	0.6	1.0
Costs for PI 3	0.01	0.02	0.032	0.05	0.65	1.0
Costs for PI 4	0.01	0.02	0.032	0.05	0.7	1.0

Table 1: Details of different problem instances (PIs) derived from synthetic datasets.

Heart Disease dataset: Each sample of the Heart Disease dataset has 12 features. We split the features into three subsets and train a logistic classifier on each subset. We associate 1st classifier with the first 7 features as input that include cholesterol readings, blood-sugar, and rest-ECG. The 2nd classifier, in addition to the 7 features, utilizes the thalach, exang, and oldpeak features; and the 3rd classifier uses all the features. For performance evaluation, the different values of costs are used in three problem instances for both real datasets, as given in Table 2. The PIMA diabetes dataset has 768 samples, whereas the Heart Disease dataset has only 297 samples. As 5000 contexts are used in our experiments, we select a sample in a round-robin fashion and give it as input to the algorithm.

Values/ Classifiers	PIMA Indian Diabetes Dataset				Heart Disease Dataset			
	Clf. 1	Clf. 2	Clf. 3	WD Fraction	Clf. 1	Clf. 2	Clf. 3	WD Fraction
Costs for PI 1	0.01	0.25	0.5	0.0692	0.01	0.25	0.5	0.1384
Costs for PI 2	0.01	0.3	0.5	0.1192	0.01	0.3	0.5	0.1454
Costs for PI 3	0.01	0.35	0.5	0.2204	0.01	0.35	0.5	0.2426

Table 2: Details of different problem instances (PIs) derived from real datasets.

³For computing disagreement probability, we assume that the feedback of any arm is independent of the feedback of other arms. Note that **USS-PD** does not need such an assumption.

Experiments Results: Through our experiments, we show that the stronger the CWD property (large value of ξ) for the problem instance, it is easier to identify the optimal arm and, hence, has lower regret, as shown in Fig. 2a. We also compare the performance of **USS-PD** with three baseline policies on problem instances derived from the Heart Disease dataset (same as the PIMA Indian Diabetes dataset). As expected, we observe that **USS-PD** outperforms the baseline policies, as shown in Fig. 2b. Note that we used $\delta = 0.05$ and $\sigma = 0.1$ in all experiments.




Figure 2: Performance of **USS-PD**.

C.1 Realizable Setting

We consider the realizable case where all contexts satisfy Eq. (8) (by fixing θ_{ij} for each (i, j) pair of arms) and WD property. Since WD holds, we can use Lemma 2 for finding the optimal arm. Note that the mean loss cannot be computed for this setting as we set parameters of disagreement probabilities instead of setting parameters for individual arms. We use an upper bound on the regret to evaluate the performance of **USS-PD** on the Synthetic dataset, as shown in Fig. 3. We repeat experiments 500 times to get a tighter confidence interval.




Figure 3: Performance of **USS-PD** for realizable setting where regret on y-axis is $\sum_{t=1}^T |C_{I_t} - C_{i_t^*}| + p_{i_t^* I_t}^{(t)}$, and it is an upper bound on the regret $\mathfrak{R}_T$ defined in Eq. (2). The value of ξ largest for Case 1, and it decreases for subsequent cases.

Regret used for Empirical Evaluation in Realizable Setting

Since the error-rate of arms is unknown, the regret defined in Eq. (2) can not be computed. Hence we

define an alternative regret, which we call pseudo regret, as follows:

$$\mathfrak{R}_T^s = \sum_{t=1}^T \left[C_{I_t} - C_{i_t^*} + p_{i_t^* I_t}^{(t)} \right].$$

It is easy to verify that the actual regret $\mathfrak{R}_T$ is upper bounded by above regret $\mathfrak{R}_T^s$ as shown follows:

$$\begin{aligned} \mathfrak{R}_T &= \sum_{t=1}^T [C_{I_t} + \gamma_{I_t}(x_t) - (C_{i_t^*} + \gamma_{i_t^*}(x_t))] \\ &= \sum_{t=1}^T [C_{I_t} - C_{i_t^*} + (\gamma_{I_t}(x_t) - \gamma_{i_t^*}(x_t))] \\ &\leq \sum_{t=1}^T [C_{I_t} - C_{i_t^*} + p_{i_t^* I_t}^{(t)}] \quad (\text{Using Lemma 1}) \\ \implies \mathfrak{R}_T &\leq \mathfrak{R}_T^s. \end{aligned}$$

C.2 Contextual Strong Dominance

We next introduce contextual strong dominance property of the problem instance.

Definition 3 (Contextual Strong Dominance (CSD) property). *A problem instance is said to satisfy CSD property if for all contexts following is true:*

$$Y^i = Y \text{ for some } i \in [K] \implies Y^j = Y, \forall j \in [K] \setminus [i].$$

We represent the set of all instances satisfies CSD property by Θ_{CSD} .

The CSD property implies that if the feedback of an arm is the same as the true reward of a given context then, the feedback of all the arms in the subsequent stages of the cascade is also the same as the true reward of a given context.

When any problem instance satisfies CSD property, the value of $\mathbb{P} \left\{ Y_t^i = Y_t, Y_t^i \neq Y_t^j | X = x_t \right\} = 0$ for $j > i$. Therefore, for any (i, j) pair of arms and context x_t the following is true:

$$\forall j > i, \gamma_i(x_t) - \gamma_j(x_t) = \mathbb{P} \left\{ Y_t^i \neq Y_t^j | X = x_t \right\}.$$

The above equation implies that CWD property holds trivially for the problem instances that satisfy CSD property as the difference of mean losses is the same as the probability of disagreement between two arms (fix arm $i = i_t^*$ for given context x_t).

C.3 Effect of adding more arms on WD property

The performance of **USS-PD** can deteriorate as we increases as the number of arms. This is because the fraction of contexts that satisfy WD property can decrease with the increase in the number of arms. To see that, consider a contextual USS problem instance with three arms where arm 1 has cost 0.1, arm 2 has cost 0.2, and arm 3 has cost 0.3. Let there be two contexts x_1 and x_2 such that classifier 2 is an optimal classifier for context x_1 and classifier 3 for the context x_2 , and both contexts satisfy WD property. When a new arm is added at the end of the classifiers cascade without changing the optimal arm for the contexts, let $p_{24}^{(1)}$ be the disagreement probability for classifier 2 and 4 for context x_1 and $p_{34}^{(2)}$ be the disagreement probability for classifier 3 and 4 for context x_2 . It is easy to verify that if cost of using classifier 4 is less than $\min\{0.2 + p_{24}^{(1)}, 0.3 + p_{34}^{(2)}\}$ then both contexts will not satisfy WD property.