
Doubly Robust Thompson Sampling with Linear Payoffs

Wonyoung Kim

Department of Statistics
Seoul National University
eraser347@snu.ac.kr

Gi-Soo Kim

Department of Industrial Engineering &
Artificial Intelligence Graduate School
UNIST
gisookim@unist.ac.kr

Myunghee Cho Paik

Department of Statistics
Seoul National University
Shepherd23 Inc.
myungheechopaik@snu.ac.kr

Abstract

A challenging aspect of the bandit problem is that a stochastic reward is observed only for the chosen arm and the rewards of other arms remain missing. The dependence of the arm choice on the past context and reward pairs compounds the complexity of regret analysis. We propose a novel multi-armed contextual bandit algorithm called Doubly Robust (DR) Thompson Sampling employing the doubly-robust estimator used in missing data literature to Thompson Sampling with contexts (LinTS). Different from previous works relying on missing data techniques (Dimakopoulou et al. [2019], Kim and Paik [2019]), the proposed algorithm is designed to allow a novel additive regret decomposition leading to an improved regret bound with the order of $\tilde{O}(\phi^{-2}\sqrt{T})$, where ϕ^2 is the minimum eigenvalue of the covariance matrix of contexts. This is the first regret bound of LinTS using ϕ^2 without the dimension of the context, d . Applying the relationship between ϕ^2 and d , the regret bound of the proposed algorithm is $\tilde{O}(d\sqrt{T})$ in many practical scenarios, improving the bound of LinTS by a factor of \sqrt{d} . A benefit of the proposed method is that it utilizes all the context data, chosen or not chosen, thus allowing to circumvent the technical definition of unsaturated arms used in theoretical analysis of LinTS. Empirical studies show the advantage of the proposed algorithm over LinTS.

1 Introduction

Contextual bandit has been popular in sequential decision tasks such as news article recommendation systems. In bandit problems, the learner sequentially pulls one arm among multiple arms and receives random rewards on each round of time. While not knowing the compensation mechanisms of rewards, the learner should make his/her decision to maximize the cumulative sum of rewards. In the course of gaining information about the compensation mechanisms through feedback, the learner should carefully balance between exploitation, pulling the best arm based on information accumulated so far, and exploration, pulling the arm that will assist in future choices, although it does not seem to be the best option at the moment. Therefore in the bandit problem, estimation or learning is an important element besides decision making.

Table 1: The shaded data are used in *complete record analysis* (left) and DR method (right) under multi-armed contextual bandit settings. The contexts, rewards and DR imputing values are denoted by X , Y , and Y^{DR} , respectively. The question mark refers to the missing reward of unchosen arms.

	$t = 1$		$t = 2$			$t = 1$		$t = 2$	
Arm 1	$X_1(1)$?	$X_1(2)$?	Arm 1	$X_1(1)$	$Y_1^{DR}(1)$	$X_1(2)$	$Y_1^{DR}(2)$
Arm 2	$X_2(1)$?	$X_{a_2}(2)$	$Y_{a_2}(2)$	Arm 2	$X_2(1)$	$Y_2^{DR}(1)$	$X_{a_2}(2)$	$Y_{a_2}^{DR}(2)$
Arm 3	$X_{a_1}(1)$	$Y_{a_1}(1)$	$X_3(2)$?	Arm 3	$X_{a_1}(1)$	$Y_{a_1}^{DR}(1)$	$X_3(2)$	$Y_3^{DR}(2)$
Arm 4	$X_4(1)$?	$X_4(2)$?	Arm 4	$X_4(1)$	$Y_4^{DR}(1)$	$X_4(2)$	$Y_4^{DR}(2)$

A challenging aspect of estimation in the bandit problem is that a stochastic reward is observed only for the chosen arm. Consequently, only the context and reward pair of the chosen arm is used for estimation, which causes dependency of the context data at the round on the past contexts and rewards. To handle this difficulty, we view bandit problems as missing data problems. The first step in handling missing data is to define full, observed, and missing data. In bandit settings, full data consist of rewards and contexts of all arms; observed data consist of full contexts for all arms and the reward for the chosen arm; missing data consist of the rewards for the arms that are not chosen. Typical estimation procedures require both rewards and contexts pairs to be observed, and the observed contexts from the unselected are discarded (see Table 1). The analysis based on the completely observed pairs only is called *complete record analysis*. Most stochastic bandit algorithms utilize estimates based on *complete record analysis*. Estimators from *complete record analysis* are known to be inefficient. In bandit setting, using the observed data whose probability of observation depends on previous rewards requires special theoretical treatment.

There are two main approaches to missing data: imputation and inverse probability weighting (IPW). Imputation is to fill in the predicted value of missing data from a specified model, and IPW is to use the observed records only but weight them by the inverse of the observation probability. The doubly robust (DR) method [Robins et al., 1994, Bang and Robins, 2005] is a combination of imputation and IPW tools. We provide a review of missing data and DR methods in supplementary materials. The robustness against model misspecification in missing data settings is insignificant in the bandit setting since the probability of observation or allocation to an arm is known. The merit of the DR method in the bandit setting is its ability to employ all the contexts including unselected arms.

We propose a novel multi-armed contextual bandit algorithm called Doubly Robust Thompson Sampling (DRTS) that applies the DR technique used in missing data literature to Thompson Sampling with linear contextual bandits (LinTS). The main thrust of DRTS is to utilize contexts information for all arms, not just chosen arms. By using the unselected, yet observed contexts, along with a novel algorithmic device, the proposed algorithm renders a unique regret decomposition which leads to a novel regret bound without resorting to the technical definition of unsaturated arms used by Agrawal and Goyal [2014]. Since categorizing the arms into saturated vs. unsaturated plays a critical role in costing extra \sqrt{d} , by circumventing it, we prove a $\tilde{O}(d\sqrt{T})$ bound of the cumulative regret in many practical occasions compared to $\tilde{O}(d^{3/2}\sqrt{T})$ shown in Agrawal and Goyal [2014].

The main contributions of this paper are as follows.

- We propose a novel contextual bandit algorithm that improves the cumulative regret bound of LinTS by a factor of \sqrt{d} (Theorem 1) in many practical scenarios (Section 4.1). This improvement is attained mainly by defining a novel set called *super-unsaturated* arms, that is utilizable due to the proposed estimator and resampling technique adopted in the algorithm.
- We provide a novel estimation error bound of the proposed estimator (Theorem 3) which depends on the minimum eigenvalue of the covariance matrix of the contexts from all arms without d .
- We develop a novel dimension-free concentration inequality for sub-Gaussian vector martingale (Lemma 4) and use it in deriving our regret bound in place of the self-normalized theorem by Abbasi-Yadkori et al. [2011].
- We develop a novel concentration inequality for the bounded matrix martingale (Lemma 6) which improves the existing result (Proposition 5) by removing the dependency on d in the bound. Lemma 6 also allows eliminating the forced sampling phases required in some bandit algorithms relying on Proposition 5 [Amani et al., 2019, Bastani and Bayati, 2020].

All missing proofs are in supplementary materials.

2 Related works

Thompson Sampling [Thompson, 1933] has been extensively studied and shown solid performances in many applications (e.g. Chapelle and Li [2011]). Agrawal and Goyal [2013] is the first to prove theoretical bounds for LinTS and an alternative proof is given by Abeille et al. [2017]. Both papers show $\tilde{O}(d^{3/2}\sqrt{T})$ regret bound, which is known as the best regret bound for LinTS. Recently, Hamidi and Bayati [2020] points out that $\tilde{O}(d^{3/2}\sqrt{T})$ could be the best possible one can get when the estimator used by LinTS is employed. In our work, we improve this regret bound by a factor of \sqrt{d} in many practical scenarios through a novel definition of super-unsaturated arms, which becomes utilizable due to the proposed estimator and resampling device implemented in the algorithm.

Our work assumes the independence of the contexts from all arms across time rounds. Some notable works have used the assumption that the contexts are independently identically distributed (IID). Leveraging the IID assumption with a margin condition, Goldenshluger and Zeevi [2013] derives a two-armed linear contextual bandit algorithm with a regret upper bound of order $O(d^3 \log T)$. Bastani and Bayati [2020] has extended this algorithm to any number of arms and improves the regret bound to $O(d^2 \log^{\frac{3}{2}} d \cdot \log T)$. The margin condition states that the gap between the expected rewards of the optimal arm and the next best arm is nonzero with some constant probability. This condition is crucial in achieving a $O(\log T)$ regret bound instead of $\tilde{O}(\sqrt{T})$. In this paper, we do not assume this margin condition, and focus on the dependence on the dimension of contexts d .

From a missing data point of view, most stochastic contextual bandit algorithms use the estimator from *complete record analysis* except Dimakopoulou et al. [2019] and Kim and Paik [2019]. Dimakopoulou et al. [2019] employs an IPW estimator that is based on the selected contexts alone. Dimakopoulou et al. [2019] proves a $\tilde{O}(d\sqrt{\epsilon^{-1}T^{1+\epsilon}N})$ regret bound for their algorithm which depends on the number of arms, N . Kim and Paik [2019] considers the high-dimensional settings with sparsity, utilizes a DR technique, and improves the regret bound in terms of the sparse dimension instead of the actual dimension of the context, d . Kim and Paik [2019] is different from ours in several aspects: the mode of exploration (ϵ -greedy vs. Thompson Sampling), the mode of regularization (Lasso vs. ridge regression); and the form of the estimator. A sharp distinction between the two estimators lies in that Kim and Paik [2019] aggregates contexts and rewards over the arms although they employ all the contexts. If we apply this aggregating estimator and DR-Lasso bandit algorithm to the low-dimensional setting, we obtain a regret bound of order $O(\frac{Nd}{\phi^2}\sqrt{T})$ when the contexts from the arms are independent. This bound is bigger than our bound by a factor of d and N . It is because the aggregated form of the estimator does not permit the novel regret decomposition derived in Section 4.2. The proposed estimator coupled with a novel algorithmic device renders the additive regret decomposition which in turn improves the order of the regret bound.

3 Proposed estimator and algorithm

3.1 Settings and assumptions

We denote a d -dimensional context for the i^{th} arm at round t by $X_i(t) \in \mathbb{R}^d$, and the corresponding random reward by $Y_i(t)$ for $i = 1, \dots, N$. We assume $\mathbb{E}[Y_i(t)|X_i(t)] = X_i(t)^T \beta$ for some unknown parameter $\beta \in \mathbb{R}^d$. At round t , the arm that the learner chooses is denoted by $a_t \in \{1, \dots, N\}$, and the optimal arm by $a_t^* := \arg \max_{i=1, \dots, N} \{X_i(t)^T \beta\}$. Let $regret(t)$ be the difference between the expected reward of the chosen arm and the optimal arm at round t , i.e., $regret(t) := X_{a_t^*}(t)^T \beta - X_{a_t}(t)^T \beta$. The goal is to minimize the sum of regrets over T rounds, $R(T) := \sum_{t=1}^T regret(t)$. The total round T is finite but possibly unknown. We also make the following assumptions.

Assumption 1. Boundedness for scale-free regrets. For all $i = 1, \dots, N$ and $t = 1, \dots, T$, we have $\|X_i(t)\|_2 \leq 1$ and $\|\beta\|_2 \leq 1$.

Assumption 2. Sub-Gaussian error. Let $\mathcal{H}_t := \bigcup_{\tau=1}^{t-1} [\{X_i(\tau)\}_{i=1}^N \cup \{a_\tau\} \cup \{Y_{a_\tau}(\tau)\}] \cup \{X_i(t)\}_{i=1}^N$ be the set of observed data at round t . For each t and i , the error $\eta_i(t) := Y_i(t) - X_i(t)^T \beta$

is conditionally zero-mean σ -sub-Gaussian for a fixed constant $\sigma \geq 0$, i.e., $\mathbb{E}[\eta_i(t)|\mathcal{H}_t] = 0$ and $\mathbb{E}[\exp(\lambda\eta_i(t))|\mathcal{H}_t] \leq \exp(\lambda^2\sigma^2/2)$, for all $\lambda \in \mathbb{R}$. Furthermore, the distribution of $\eta_i(t)$ does not depend on the choice at round t , i.e. a_t .

Assumption 3. Independently distributed contexts. The stacked contexts vectors $\{X_i(1)\}_{i=1}^N, \dots, \{X_i(T)\}_{i=1}^N \in \mathbb{R}^{dN}$ are independently distributed.

Assumption 4. Positive minimum eigenvalue of the average of covariance matrices. For each t , there exists a constant $\phi^2 > 0$ such that $\lambda_{\min}\left(\mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N X_i(t)X_i(t)^T\right]\right) \geq \phi^2$.

Assumptions 1 and 2 are standard in stochastic bandit literature Agrawal and Goyal [2013]. We point out that given round t , Assumption 3 allows that the contexts among different arms, $X_1(t), \dots, X_N(t)$ are correlated to each other. Assumption 3 is weaker than the assumption of IID, and the IID condition is considered by Goldenshluger and Zeevi [2013] and Bastani and Bayati [2020]. As Bastani and Bayati [2020] points out, the IID assumption is reasonable in some practical settings, including clinical trials, where health outcomes of patients are independent of those of other patients. Both Goldenshluger and Zeevi [2013] and Bastani and Bayati [2020] address the problem where the contexts are equal across all arms, i.e. $X(t) = X_1(t) = \dots = X_N(t)$, while our work admits different contexts over all arms. Assumption 4 guarantees that the average of covariance matrices of contexts over the arms is well-behaved so that the inverse of the sample covariance matrix is bounded by the spectral norm. This assumption helps controlling the estimation error of β in linear regression models. Similar assumptions are adopted in existing works in the bandit setting [Goldenshluger and Zeevi, 2013, Amani et al., 2019, Li et al., 2017, Bastani and Bayati, 2020].

3.2 Doubly robust estimator

To describe the contextual bandit DR estimator, let $\pi_i(t) := \mathbb{P}(a_t = i | \mathcal{H}_t) > 0$ be the probability of selecting arm i at round t . We define a DR pseudo-reward as

$$Y_i^{DR}(t) = \left\{1 - \frac{\mathbb{I}(i = a_t)}{\pi_i(t)}\right\} X_i(t)^T \check{\beta}_t + \frac{\mathbb{I}(i = a_t)}{\pi_i(t)} Y_{a_t}(t), \quad (1)$$

for some $\check{\beta}_t$ depending on \mathcal{H}_t . Background of missing data methods and derivation of the DR pseudo-reward is provided in the supplementary material. Now, we propose our new estimator $\hat{\beta}_t$ with a regularization parameter λ_t as below:

$$\hat{\beta}_t = \left(\sum_{\tau=1}^t \sum_{i=1}^N X_i(\tau)X_i(\tau)^T + \lambda_t I\right)^{-1} \left(\sum_{\tau=1}^t \sum_{i=1}^N X_i(\tau)Y_i^{DR}(\tau)\right). \quad (2)$$

Harnessing the pseudo-rewards defined in (1), we can make use of all contexts rather than just selected contexts. The DR estimator by Kim and Paik [2019] utilizes all contexts but has a different form from ours. While Kim and Paik [2019] uses Lasso estimator with pseudo-rewards *aggregated* over all arms, we use ridge regression estimator with pseudo-rewards in (1) which are defined *separately* for each $i = 1, \dots, N$. This seemingly small but important difference in forms paves a way in rendering our unique regret decomposition and improving the regret bound.

3.3 Algorithm

In this subsection, we describe our proposed algorithm, DRTS which adapts DR technique to LinTS. The DRTS is presented in Algorithm 1. Distinctive features of DRTS compared to LinTS include the novel estimator and the resampling technique. At each round $t \geq 1$, the algorithm samples $\tilde{\beta}_i(t)$ from the distribution $N(\hat{\beta}_{t-1}, v^2 V_{t-1}^{-1})$ for each i independently. Let $\tilde{Y}_i(t) := X_i(t)^T \tilde{\beta}_i(t)$ and $m_t := \arg \max_i \tilde{Y}_i(t)$. We set m_t as a candidate action and compute $\tilde{\pi}_{m_t}(t) := \mathbb{P}(\tilde{Y}_{m_t}(t) = \max_i \tilde{Y}_i(t) | \mathcal{H}_t)$.¹ If $\tilde{\pi}_{m_t}(t) > \gamma$, then the arm m_t is selected, i.e., $a_t = m_t$. Otherwise, the algorithm resamples $\tilde{\beta}_i(t)$ until it finds another arm satisfying $\tilde{\pi}_i(t) > \gamma$ up to a predetermined fixed value M_t . Section A.3 in supplementary materials describes issues related to M_t including a suitable choice of M_t .

¹This computation is known to be challenging but employing the independence among $\tilde{\beta}_1(t), \dots, \tilde{\beta}_N(t)$, we derive an explicit form approximating $\tilde{\pi}_{m_t}(t)$ in supplementary materials Section H.1.

Algorithm 1 Doubly Robust Thompson Sampling for Linear Contextual Bandits (DRTS)

Input: Exploration parameter $v > 0$, Regularization parameter $\lambda > 0$, Selection probability threshold $\gamma \in [1/(N+1), 1/N]$, Imputation estimator $\check{\beta}_u = f(\{X(\tau), Y_{a_\tau}(\tau)\}_{\tau=1}^{u-1})$, Number of maximum possible resampling M_t .

Set $F_0 = 0$, $W_0 = 0$, $\hat{\beta}_0 = 0$ and $V_0 = \lambda I$

for $t = 1$ **to** T **do**

Observe contexts $\{X_i(t)\}_{i=1}^N$.

Sample $\tilde{\beta}_1(t), \dots, \tilde{\beta}_N(t)$ from $N(\hat{\beta}_{t-1}, v^2 V_{t-1}^{-1})$ independently. Compute $\tilde{Y}_i(t) = X_i(t)^T \tilde{\beta}_i(t)$

Observe a candidate action $m_t := \arg \max_i \tilde{Y}_i(t)$.

Compute $\tilde{\pi}_{m_t}(t) := \mathbb{P} \left(\max_i \tilde{Y}_i(t) = \tilde{Y}_{m_t}(t) \mid \mathcal{H}_t \right)$.

for $l = 1$ **to** M_t **do**

if $\tilde{\pi}_{m_t}(t) \leq \gamma$ **then**

Sample another $\tilde{\beta}_1(t), \dots, \tilde{\beta}_N(t)$, observe another m_t , and update $\tilde{\pi}_{m_t}(t)$.

else

Break.

end if

end for

Set $a_t = m_t$, and play arm a_t .

Observe reward $Y_{a_t}(t)$ and compute $Y_t^{DR}(t)$

$F_t = F_{t-1} + \sum_{i=1}^N X_i(t) Y_i^{DR}(t)$; $W_t = W_{t-1} + \sum_{i=1}^N X_i(t) X_i(t)^T$; $V_t = W_t + \lambda \sqrt{t} I$

$\hat{\beta}_t = V_t^{-1} F_t$

Update $\check{\beta}_{t+1}$ for next round.

end for

The resampling step is incorporated to avoid small values of the probability of selection so that the pseudo-reward in (1) is numerically stable. A naive remedy to stabilize the pseudo-reward is to use $\max\{\pi_i(t), \gamma\}$, which fails to leading to our regret bound since it induces bias and also cannot guarantee that the selected arm is in the super-unsaturated arms defined in (5) with high probability (For details, see Section 4.2). The resampling step implemented in the proposed algorithm is designed to solve these problems.

4 Theoretical results

Our theoretical results are organized as follows. In Section 4.1, we provide the main result, the cumulative regret bound of $\tilde{O}(\phi^{-2} \sqrt{T})$ of DRTS. The main thrust of deriving the regret bound is to define super-unsaturated arms. In Section 4.2 we introduce the definition of super-unsaturated arms and show how it admits a novel decomposition of the regret into two additive terms as in (6). In Section 4.3 we bound each term of the decomposed regret bounds (6). The first term is the estimation error, and Theorem 3 finds its bound. In the course of proving Theorem 3, we need Lemma 4, which plays a similar role to the self-normalized theorem of Abbasi-Yadkori et al. [2011]. We conclude the section by presenting Lemma 6 and bound the second term of (6).

4.1 An improved regret bound

Theorem 1 provides the regret bound of DRTS in terms of the minimum eigenvalue without d .

Theorem 1. *Suppose that Assumptions 1-4 hold. If $\check{\beta}_t$ in Algorithm 1 satisfies $\|\check{\beta}_t - \beta\|_2 \leq b$ for a constant $b > 0$, for all $t = 1, \dots, T$, then with probability $1 - 2\delta$, the cumulative regret by time T for DRTS algorithm is bounded by*

$$R(T) \leq 2 + \frac{4C_{b,\sigma}}{\phi^2} \sqrt{T \log \frac{12T^2}{\delta}} + \frac{2\sqrt{2T}}{\phi\sqrt{N}}, \quad (3)$$

where $C_{b,\sigma}$ is a constant which depends only on b and σ .

The bound (3) has a rate of $O(\phi^{-2}\sqrt{T})$. The relationship between the dimension d and the minimum eigenvalue ϕ^2 can be shown by

$$d\phi^2 = \frac{d}{N}\lambda_{\min}\left(\mathbb{E}\sum_{i=1}^N X_i(t)X_i(t)^T\right) \leq \frac{1}{N}\mathbb{E}\sum_{i=1}^N \text{Tr}(X_i(t)X_i(t)^T) = \frac{1}{N}\mathbb{E}\sum_{i=1}^N \|X_i(t)\|_2^2 \leq 1.$$

This implies $\phi^{-2} \geq d$,² but there are many practical scenarios such that $\phi^{-2} = O(d)$ holds. Bastani et al. [2021] identifies such examples including the uniform distribution and truncated multivariate normal distributions. When the context has uniform distribution on the unit ball, $\phi^{-2} = d + 2$. When the context has truncated multivariate normal distribution with mean 0 and covariance Σ , we can set $\phi^{-2} = (d + 2)\exp(\frac{1}{2\lambda_{\min}(\Sigma)})$. For more examples, we refer to Bastani et al. [2021]. Furthermore, regardless of distributions, $\phi^{-2} = O(d)$ holds when the correlation structure has the row sum of off-diagonals independent of the dimension, for example, AR(1), tri-diagonal, block-diagonal matrices. In these scenarios, the regret bound in (3) becomes $\tilde{O}(d\sqrt{T})$. Compared to the previous bound of LinTS [Agrawal and Goyal, 2014, Abeille et al., 2017], we obtain a better regret bound by the factor of \sqrt{d} for identified practical cases.

As for the imputation estimator $\tilde{\beta}_t$, we assume that $\|\tilde{\beta}_t - \beta\|_2 \leq b$, where b is an absolute constant. We suggest two cases which guarantee this assumption. First, if a biased estimator is used, we can rescale the estimator so that its l_2 -norm is bounded by some constant $C > 0$. Then, $\|\tilde{\beta}_t - \beta\|_2 \leq \|\tilde{\beta}_t\|_2 + \|\beta\|_2 \leq C + 1$ and $b = C + 1$. Second, consistent estimators such as ridge estimator or the least squared estimator satisfy the condition since $\|\tilde{\beta}_t - \beta\|_2 = O(d\sqrt{\log t/t})$. The term d is cancelled out when $t \geq t_d$, where t_d is the minimum integer that satisfies $\log t/t \leq d^{-2}$. In these two cases, we can find a constant b which satisfies the assumption on the imputation estimator $\tilde{\beta}_t$.

4.2 Super-unsaturated arms and a novel regret decomposition

The key element in deriving (3) is to decompose the regret into two additive terms as in (6). To allow such decomposition to be utilizable, we need to define a novel set of arms called super-unsaturated arms, which replaces the role of unsaturated arms in [Agrawal and Goyal, 2014]. The super-unsaturated arms are formulated so that the chosen arm is included in this set with high probability. For each i and t , let $\Delta_i(t) := X_{a_i^*}(t)^T\beta - X_i(t)^T\beta$. Define $A_t := \sum_{\tau=1}^t X_{a_\tau}(\tau)X_{a_\tau}(\tau)^T + \lambda I$ and $V_t := \sum_{\tau=1}^t \sum_{i=1}^N X_i(\tau)X_i(\tau)^T + \lambda_t I$. For the sake of contrast, recall the definition of unsaturated arms by Agrawal and Goyal [2014]

$$U_t := \left\{i : \Delta_i(t) \leq g_t \|X_i(t)\|_{A_{t-1}^{-1}}\right\}, \quad (4)$$

where $g_t := C\sqrt{d\log(t/\delta)}\min\{\sqrt{d}, \sqrt{\log N}\}$ for some constant $C > 0$. This g_t is constructed to ensure that there exists a positive lower bound for the probability that the selected arm is unsaturated. In place of (4), we define a set of super-unsaturated arms for each round t by

$$N_t := \left\{i : \Delta_i(t) \leq 2\left\|\hat{\beta}_{t-1} - \beta\right\|_2 + \sqrt{\|X_{a_i^*}(t)\|_{V_{t-1}^{-1}}^2 + \|X_i(t)\|_{V_{t-1}^{-1}}^2}\right\}. \quad (5)$$

While $g_t \|X_i(t)\|_{A_{t-1}^{-1}}$ in (4) is normalized with only selected contexts, the second term in the right hand side of (5) is normalized with all contexts including $X_{a_i^*}(t)$, the contexts of the optimal arm. This bound of $\Delta_i(t)$ plays a crucial role in bounding the regret with a novel decomposition as in (6). The following Lemma shows a lower bound of the probability that the candidate arm is super-unsaturated.

Lemma 2. *For each t , let $m_t := \arg \max_i \tilde{Y}_i(t)$ and let N_t be the super-unsaturated arms defined in (5). For any given $\gamma \in [1/(N+1), 1/N]$, set $v = (2\log(N/(1-\gamma N)))^{-1/2}$. Then, $\mathbb{P}(m_t \in N_t | \mathcal{H}_t) \geq 1 - \gamma$.*

Lemma 2 directly contributes to the reduction of \sqrt{d} in the hyperparameter v . In Agrawal and Goyal [2014], to prove a lower bound of $\mathbb{P}(a_t \in U_t | \mathcal{H}_t)$, it is required to set $v = \sqrt{9d\log(t/\delta)}$, with the

²Some previous works assume $\phi^{-2} = O(1)$ even when $\|X_i(t)\|_2 \leq 1$ (e.g. Li et al. [2017]). As pointed out by Ding et al. [2021], this assumption is unrealistic and the reported regret bound should be multiplied by $O(d)$.

order of \sqrt{d} . In contrast, Lemma 2 shows that v does not need to depend on d due to the definition of super-unsaturated arms in (5). In this way, we obtain a lower bound of $\mathbb{P}(m_t \in N_t | \mathcal{H}_t)$ without costing extra \sqrt{d} .

Using the lower bound, we can show that the resampling scheme allows the algorithm to choose the super-unsaturated arms with high probability. For all $i \notin N_t$,

$$\tilde{\pi}_i(t) := \mathbb{P}(m_t = i | \mathcal{H}_t) \leq \mathbb{P}(\cup_{j \notin N_t} \{m_t = j\} | \mathcal{H}_t) = \mathbb{P}(m_t \notin N_t | \mathcal{H}_t) \leq \gamma,$$

where the last inequality holds due to Lemma 2. Thus, in turn, if $\tilde{\pi}_i(t) > \gamma$, then $i \in N_t$. This means that $\{i : \tilde{\pi}_i(t) > \gamma\}$ is a subset of N_t and

$$\{a_t \in \{i : \tilde{\pi}_i(t) > \gamma\}\} \subset \{a_t \in N_t\}.$$

Hence, the probability of the event $\{a_t \in N_t\}$ is greater than the probability of sampling any arm which satisfies $\tilde{\pi}_i(t) > \gamma$. Therefore, with resampling, the event $\{a_t \in N_t\}$ occurs with high probability. (See supplementary materials Section A for details.)

When the algorithm chooses the arm from the super-unsaturated set, i.e., when $a_t \in N_t$ happens, (5) implies

$$\Delta_{a_t}(t) \leq 2 \left\| \hat{\beta}_{t-1} - \beta \right\|_2 + \sqrt{\|X_{a_t}(t)\|_{V_{t-1}^{-1}}^2 + \|X_{a_t}(t)\|_{V_{t-1}}^2}. \quad (6)$$

By definition, $\Delta_{a_t}(t) = \text{regret}(t)$ and the regret at round t can be expressed as the two additive terms, which presents a stark contrast with multiplicative decomposition of the regret in Agrawal and Goyal [2014]. In section 4.3 we show how each term can be bounded with separate rate.

4.3 Bounds for the cumulative regret

We first bound the leading term of (6) and introduce a novel estimation error bound free of d for the contextual bandit DR estimator.

Theorem 3. (A dimension-free estimation error bound for the contextual bandit DR estimator.) *Suppose Assumptions 1-4 hold. For each $t = 1, \dots, T$, let $\check{\beta}_t$ be any \mathcal{H}_t -measurable estimator satisfying $\|\check{\beta}_t - \beta\|_2 \leq b$, for some constant $b > 0$. For each i and t , assume that $\pi_i(t) > 0$ and that there exists $\gamma \in [1/(N+1), 1/N]$ such that $\pi_{a_t}(t) > \gamma$. Given any $\delta \in (0, 1)$, set $\lambda_t = 4\sqrt{2}N\sqrt{t \log \frac{12t^2}{\delta}}$. Then with probability at least $1 - \delta$, the estimator $\hat{\beta}_t$ in (2) satisfies*

$$\left\| \hat{\beta}_t - \beta \right\|_2 \leq \frac{C_{b,\sigma}}{\phi^2 \sqrt{t}} \sqrt{\log \frac{12t^2}{\delta}}, \quad (7)$$

for all $t = 1, \dots, T$, where the constant $C_{b,\sigma}$ which depends only on b and σ .

In bandit literature, estimation error bounds typically include a term involving d which emerges from using the following two Lemmas: (i) the self-normalized bound for vector-valued martingales [Abbasi-Yadkori et al., 2011, Theorem 1], and (ii) the concentration inequality for the covariance matrix [Tropp, 2015, Corollary 5.2]. Instead of using (i) and (ii), we develop the two dimension-free bounds in Lemmas 4 and 6, to replace (i) and (ii), respectively. With the two Lemmas, we eliminate the dependence on d and express the estimation error bound with ϕ^2 alone.

Lemma 4. (A dimension-free bound for vector-valued martingales.) *Let $\{\mathcal{F}_\tau\}_{\tau=1}^t$ be a filtration and $\{\eta(\tau)\}_{\tau=1}^t$ be a real-valued stochastic process such that $\eta(\tau)$ is \mathcal{F}_τ -measurable. Let $\{X(\tau)\}_{\tau=1}^t$ be an \mathbb{R}^d -valued stochastic process where $X(\tau)$ is $\mathcal{F}_{\tau-1}$ -measurable and $\|X(\tau)\|_2 \leq 1$. Assume that $\{\eta(\tau)\}_{\tau=1}^t$ are σ -sub-Gaussian as in Assumption 2. Then with probability at least $1 - \delta/t^2$, there exists an absolute constant $C > 0$ such that*

$$\left\| \sum_{\tau=1}^t \eta(\tau) X(\tau) \right\|_2 \leq C\sigma\sqrt{t} \sqrt{\log \frac{4t^2}{\delta}}. \quad (8)$$

Compared to Theorem 1 of Abbasi-Yadkori et al. [2011], our bound (8) does not involve d , yielding a dimension-free bound for vector-valued martingales. However, the bound (8) has \sqrt{t} term which comes from using $\|\cdot\|_2$ instead of the self-normalized norm $\|\cdot\|_{V_t^{-1}}$.

To complete the proof of Theorem 3, we need the following condition,

$$\lambda_{\min}(V_t) \geq ct, \quad (9)$$

for some constant $c > 0$. Li et al. [2017] points out that satisfying (9) is challenging. To overcome this difficulty, Amani et al. [2019] and Bastani and Bayati [2020] use an assumption on the covariance matrix of contexts and a concentration inequality for matrix to prove (9), described as follows.

Proposition 5. [Tropp, 2015, Theorem 5.1.1] *Let $P(1), \dots, P(t) \in \mathbb{R}^{d \times d}$ be the symmetric matrices such that $\lambda_{\min}(P(\tau)) \geq 0$, $\lambda_{\max}(P(\tau)) \leq L$ and $\lambda_{\min}(\mathbb{E}[P(\tau)]) \geq \phi^2$, for all $\tau = 1, 2, \dots, t$. Then,*

$$\mathbb{P} \left(\lambda_{\min} \left(\sum_{\tau=1}^t P(\tau) \right) \leq \frac{t\phi^2}{2} \right) \leq d \exp \left(-\frac{t\phi^2}{8L} \right). \quad (10)$$

To prove (9) using (10) with probability at least $1 - \delta$, for $\delta \in (0, 1)$, it requires $t \geq \frac{8L}{\phi^2} \log \frac{d}{\delta}$. Thus, one can use (10) only after $O(\phi^{-2} \log d)$ rounds. Due to this requirement, Bastani and Bayati [2020] implements the forced sampling techniques for $O(N^2 d^4 (\log d)^2)$ rounds, and Amani et al. [2019] forces to select arms randomly for $O(\phi^{-2} \log d)$ rounds. These mandatory exploration phase empirically prevents the algorithm choosing the optimal arm. An alternative form of matrix Chernoff inequality for adapted sequences is Theorem 3 in Tropp [2011], but the bound also has a multiplicative factor of d . Instead of applying Proposition 5 to prove (9), we utilize a novel dimension-free concentration inequality stated in the following Lemma.

Lemma 6. (A dimension-free concentration bound for symmetric bounded matrices.) *Let $\|A\|_F$ be a Frobenius norm of a matrix A . Let $\{P(\tau)\}_{\tau=1}^t \in \mathbb{R}^{d \times d}$ be the symmetric matrices adapted to a filtration $\{\mathcal{F}_\tau\}_{\tau=1}^t$. For each $\tau = 1, \dots, t$, suppose that $\|P(\tau)\|_F \leq c$, for some $c > 0$ and $\lambda_{\min}(\mathbb{E}[P(\tau)|\mathcal{F}_{\tau-1}]) \geq \phi^2 > 0$, almost surely. For given any $\delta \in (0, 1)$, set $\lambda_t \geq 4\sqrt{2}c\sqrt{t} \sqrt{\log \frac{4t^2}{\delta}}$. Then with probability at least $1 - \delta/t^2$,*

$$\lambda_{\min} \left(\sum_{\tau=1}^t P(\tau) + \lambda_t I \right) \geq \phi^2 t. \quad (11)$$

Lemma 6 shows that setting λ_t with \sqrt{t} rate guarantees (9) for all $t \geq 1$. We incorporate λ_t stated in Lemma 6 in our estimator (2), and show in Section 5 that the DR estimator regularized with λ_t outperforms estimators from other contextual bandit algorithms in early rounds.

We obtain the bounds free of d in Lemmas 4 and 6 mainly by applying Lemma 2.3 in Lee et al. [2016] which states that any Hilbert space martingale can be reduced to \mathbb{R}^2 . Thus, we can project the vector-valued (or the matrix) martingales to \mathbb{R}^2 -martingales, and reduce the dimension from d (or d^2) to 2. Then we apply Azuma-Hoeffding inequality just twice, instead of d times. In this way, Lemma 6 provides a novel dimension-free bound for the covariance matrix.

Lemmas 4 and 6 can be applied to other works to improve the existing bounds. For example, using these Lemmas, the estimation error bound of Bastani and Bayati [2020] can be improved by a factor of $\log d$. Proposition EC.1 of Bastani and Bayati [2020] provides an estimation error bound for the ordinary least square estimator by using Proposition 5 and bounding all values of d coordinates. By applying Lemmas 4 and 6, one does not have to deal with each coordinate and eliminate dependence on d .

Using Lemma 6, we can bound the second term of the regret in (6) as follows. For $j = 1, \dots, N$

$$\|X_j(t)\|_{V_{t-1}^{-1}} \leq \|X_j(t)\|_2 \sqrt{\|V_{t-1}^{-1}\|_2} \leq \lambda_{\min}(V_{t-1})^{-1/2} \leq \frac{1}{\sqrt{\phi^2 N(t-1)}}. \quad (12)$$

Finally, we are ready to bound $\text{regret}(t)$ in (6).

Lemma 7. *Suppose the assumptions in Theorem 1 hold. Then with probability at least $1 - 2\delta$,*

$$\text{regret}(t) \leq \frac{2C_{b,\sigma}}{\phi^2 \sqrt{t-1}} \sqrt{\log \frac{12t^2}{\delta}} + \frac{\sqrt{2}}{\phi \sqrt{N(t-1)}}, \quad (13)$$

for all $t = 2, \dots, T$.

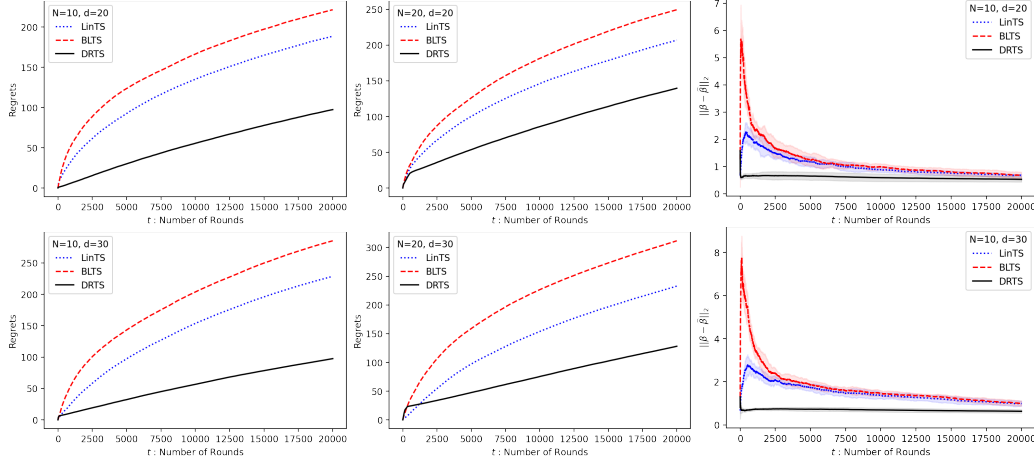


Figure 1: A Comparison of cumulative regrets and estimation errors of LinTS, BLTS and DRTS. Each line shows the averaged cumulative regrets (estimation errors, resp.) and the shaded area in the right two figures represents the standard deviations over 10 repeated experiments.

Proof. Since a_t is shown to be super-unsaturated with high probability, we can use (6) to have $\text{regret}(t) \leq 2\|\hat{\beta}_{t-1} - \beta\|_2 + \sqrt{\|X_{a_t^*}(t)\|_{V_{t-1}}^2 + \|X_{a_t}(t)\|_{V_{t-1}}^2}$, for all $t = 2, \dots, T$. We see that the first term is bounded by Theorem 3, and the second term by (12). Note that to prove Theorem 1, Lemma 6 is invoked, and the event (11) of Lemma 6 is a subset of that in (7). Therefore (13) holds with probability at least $1 - 2\delta$ instead of $1 - 3\delta$. Details are given in supplementary materials. \square

Lemma 7 shows that the regret at round t does not exceed a $O(\phi^{-2}t^{-1/2})$ bound when $a_t \in N_t$, which is guaranteed in our algorithm via resampling with high probability (See Section A.3 for details). This concludes the proof of Theorem 1.

5 Simulation studies

In this section, we compare the performances of the three algorithms: (i) LinTS [Agrawal and Goyal, 2013], (ii) BLTS [Dimakopoulou et al., 2019], and (iii) the proposed DRTS. We use simulated data described as follows. The number of arms N is set to 10 or 20, and the dimension of contexts d is set to 20 or 30. For each element of the contexts $j = 1, \dots, d$, we generate $[X_{1j}(t), \dots, X_{Nj}(t)]$ from a normal distribution $\mathcal{N}(\mu_N, V_N)$ with mean $\mu_{10} = [-10, -8, \dots, -2, 2, \dots, 8, 10]^T$, or $\mu_{20} = [-20, -18, \dots, -2, 2, \dots, 18, 20]^T$, and the covariance matrix $V_N \in \mathbb{R}^{N \times N}$ has $V_N(i, i) = 1$ for every i and $V_N(i, k) = \rho$ for every $i \neq k$. We set $\rho = 0.5$ and truncate the sampled contexts to satisfy $\|X_i(t)\|_2 \leq 1$. To generate the stochastic rewards, we sample $\eta_i(t)$ independently from $\mathcal{N}(0, 1)$. Each element of β follows a uniform distribution, $\mathcal{U}(-1/\sqrt{d}, 1/\sqrt{d})$.

All three algorithms have v as an input parameter which controls the variance of $\tilde{\beta}_i(t)$. BLTS and DRTS require a positive threshold γ which truncates the selection probability. We consider $v \in \{0.001, 0.01, 0.1, 1\}$ in all three algorithms, $\gamma \in \{0.01, 0.05, 0.1\}$ for BLTS, and set $\gamma = 1/(N + 1)$ in DRTS. Then we report the minimum regrets among all combinations. The regularization parameter is $\lambda_t = \sqrt{t}$ in DRTS and $\lambda_t = 1$ in both LinTS and BLTS. To obtain an imputation estimator $\tilde{\beta}_t$ required in DRTS, we use ridge regression with $\{X_{a_\tau}(\tau), Y_{a_\tau}(\tau)\}_{\tau=1}^{t-1}$, for each round t . Other implementation details are in supplementary materials.

Figure 1 shows the average of the cumulative regrets and the estimation error $\|\hat{\beta}_t - \beta\|_2$ of the three algorithms based on 10 replications. The figures in the two left columns show the average cumulative regret according to the number of rounds with the best set of hyperparameters for each algorithm. The total rounds are $T = 20000$. The figures in the third columns show the average of the estimation error $\|\hat{\beta}_t - \beta\|_2$. In the early stage, the estimation errors of LinTS and BLTS increase rapidly, while that of DRTS is stable. The stability of the DR estimator follows possibly by using full contexts and

the regularization parameter $\lambda_t = \sqrt{t}$. This yields a large margin of estimation error among LinTS, BLTS and DRTS, especially when the dimension is large.

6 Conclusion

In this paper, we propose a novel algorithm for stochastic contextual linear bandits. Viewing the bandit problem as a missing data problem, we use the DR technique to employ all contexts including those that are not chosen. With the definition of super-unsaturated arms, we show a regret bound which only depends on the minimum eigenvalue of the sample covariance matrices. This new bound has $\tilde{O}(d\sqrt{T})$ rate in many practical scenarios, which is improved by a factor of \sqrt{d} compared to the previous LinTS regret bounds. Simulation studies show that the proposed algorithm performs better than other LinTS algorithms in a large dimension.

Acknowledgements

This work is supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT, No.2020R1A2C1A01011950) (Wonyoung Kim and Myunghee Cho Paik), and by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2020-0-01336, Artificial Intelligence Graduate School Program(UNIST)) and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT, No.2021R1G1A100980111) (Gi-Soo Kim). Wonyoung Kim was also supported by Hyundai Chung Mong-koo foundation.

References

- 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.
- Marc Abeille, Alessandro Lazaric, et al. Linear thompson sampling revisited. *Electronic Journal of Statistics*, 11(2):5165–5197, 2017.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs, 2014.
- Sanae Amani, Mahnoosh Alizadeh, and Christos Thrampoulidis. Linear stochastic bandits under safety constraints. In *Advances in Neural Information Processing Systems*, pages 9252–9262, 2019.
- Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal, Second Series*, 19(3):357–367, 1967.
- Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.
- Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294, 2020.
- Hamsa Bastani, Mohsen Bayati, and Khashayar Khosravi. Mostly exploration-free algorithms for contextual bandits. *Management Science*, 67(3):1329–1349, 2021.
- Olivier Chapelle and Lihong Li. An empirical evaluation of thompson sampling. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2249–2257. Curran Associates, Inc., 2011.
- Fan Chung and Linyuan Lu. Concentration inequalities and martingale inequalities: a survey. *Internet Mathematics*, 3(1):79–127, 2006.

- Maria Dimakopoulou, Zhengyuan Zhou, Susan Athey, and Guido Imbens. Balanced linear contextual bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3445–3453, 2019.
- Qin Ding, Cho-Jui Hsieh, and James Sharpnack. An efficient algorithm for generalized linear bandit: Online stochastic gradient descent and thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 1585–1593. PMLR, 2021.
- Alexander Goldenshluger and Assaf Zeevi. A linear response bandit problem. *Stochastic Systems*, 3(1):230–261, 2013.
- Nima Hamidi and Mohsen Bayati. On worst-case regret of linear thompson sampling, 2020.
- Gi-Soo Kim and Myunghee Cho Paik. Doubly-robust lasso bandit. In *Advances in Neural Information Processing Systems*, pages 5869–5879, 2019.
- James R. Lee, Yuval Peres, and Charles K. Smart. A gaussian upper bound for martingale small-ball probabilities. *Ann. Probab.*, 44(6):4184–4197, 11 2016.
- Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2071–2080. JMLR. org, 2017.
- James M. Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.
- William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Joel A Tropp. User-friendly tail bounds for matrix martingales. Technical report, CALIFORNIA INST OF TECH PASADENA, 2011.
- Joel A Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See supplementary materials.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] Since this work focuses on the theoretical regret bound of an algorithm, the societal impacts of the results are too indirect and too broad to discuss.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 3.1
 - (b) Did you include complete proofs of all theoretical results? [Yes] See supplementary materials
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See supplementary materials
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 5
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Only in the estimation error results in Figure 1 for good visualization.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No] The computation in our experiments is too light to characterize and our work focuses on the theoretical analysis of our algorithm.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Detailed analysis of the resampling

In this section, we give details about the issues which can be raised from the resampling in our algorithm.

A.1 Precise definition of action selection

We give precise definition of the action at round t , a_t . For each round $t \geq 2$, given \mathcal{H}_t , let $a_t^{(1)}, a_t^{(2)}, \dots, a_t^{(M_t)}$ to be maximum possible sequence of actions to be resampled. These actions are IID, with $\mathbb{P}(a_t^{(1)} = i | \mathcal{H}_t) = \tilde{\pi}_i(t)$ for $i = 1, \dots, N$. Define a subset of arms $\tilde{\Gamma}_t := \{i : \tilde{\pi}_i(t) > \gamma\}$ and a stopping time

$$\mathcal{T} := \inf\{m \geq 1 : a_t^{(m)} \in \tilde{\Gamma}_t\} \quad (14)$$

with respect to the filtration $\mathcal{F}_m := \mathcal{H}_t \cup \{a_t^{(1)}, \dots, a_t^{(m)}\}$. Since the algorithm stops resampling when the candidate action is in $\tilde{\Gamma}_t$, the stopping time \mathcal{T} is the actual number of resampling in algorithm. Thus we can write the action after resampling as $a_t := a_t^{(\min\{\mathcal{T}, M_t\})}$.

A.2 Computing the probability of selection

The probability of selection $\pi_i(t) := \mathbb{P}(a_t = i | \mathcal{H}_t)$ is not the same as $\tilde{\pi}_i(t)$ due to resampling. This might cause the problem of computing $\pi_i(t)$ which is essential to compute $Y_i^{DR}(t)$. However, with the precise definition of a_t , we can derive a closed form for $\pi_i(t)$.

First, we consider two cases separately: (i) the case when the resampling succeeds and (ii) the case when the resampling fails and the maximum possible number of resampling runs out. In case (i), $a_t \in \tilde{\Gamma}_t$, and for any $i \in \tilde{\Gamma}_t$, we have

$$\begin{aligned} \mathbb{P}(a_t = i | \mathcal{H}_t) &= \mathbb{P}(\mathcal{T} \leq M_t, a_t^{(\mathcal{T})} = i | \mathcal{H}_t) \\ &= \sum_{m=1}^{M_t} \mathbb{P}(\mathcal{T} = m, a_t^{(m)} = i | \mathcal{H}_t) \\ &= \sum_{m=1}^{M_t} \mathbb{P}(a_t^{(m)} = i | \mathcal{H}_t) \left(\prod_{j=0}^{m-1} \mathbb{P}(a_t^{(j)} \notin \tilde{\Gamma}_t | \mathcal{H}_t) \right) \\ &= \tilde{\pi}_i(t) \sum_{m=1}^{M_t} \left(1 - \sum_{i \in \tilde{\Gamma}_t} \tilde{\pi}_i(t) \right)^{m-1} \\ &= \tilde{\pi}_i(t) \frac{1 - (1 - \sum_{i \in \tilde{\Gamma}_t} \tilde{\pi}_i(t))^{M_t}}{\sum_{i \in \tilde{\Gamma}_t} \tilde{\pi}_i(t)}. \end{aligned} \quad (15)$$

Now, for the case (ii) $a_t \notin \tilde{\Gamma}_t$, and for any $i \notin \tilde{\Gamma}_t$, we have

$$\begin{aligned} \mathbb{P}(a_t = i | \mathcal{H}_t) &= \mathbb{P}(\mathcal{T} > M_t, a_t^{(M_t)} = i | \mathcal{H}_t) \\ &= \mathbb{P}\left(\bigcap_{m=1}^{M_t-1} \{a_t^{(m)} \notin \tilde{\Gamma}_t\}, a_t^{(M_t)} = i \mid \mathcal{H}_t\right) \\ &= \left(1 - \sum_{i \in \tilde{\Gamma}_t} \tilde{\pi}_i(t)\right)^{M_t-1} \tilde{\pi}_i(t). \end{aligned} \quad (16)$$

With (15) and (16), we can compute $\pi_i(t)$ for all $i = 1, \dots, N$.

A.3 The number of maximum possible resampling

The proposed algorithm attempts resampling up to M_t times to find an arm in $\{i : \tilde{\pi}_i(t) > \gamma\}$. The main point in selecting M_t is to bound the probability that the resampling fails in finding an arm whose selection probability exceeds γ for some δ , i.e.,

$$\mathbb{P}(a_t \notin \{i : \tilde{\pi}_i(t) > \gamma\}) \leq \delta/t^2. \quad (17)$$

Intuitively, as M_t increases, we have more opportunities for resampling and the probability that the resampling fails in finding arms in $\{i : \tilde{\pi}_i(t) > \gamma\}$ decreases. Since $\gamma < 1/N$, there exists j such that $\tilde{\pi}_j(t) > \gamma$, and the probability that the resampling fails is less than $1 - \gamma$ in each resampling trial.

Specifically, we can achieve (17) by choosing M_t as a minimum integer that exceeds $\log \frac{t^2}{\delta} / \log \frac{1}{1-\gamma}$. For any given $\delta \in (0, 1)$, the event $\{a_t \in \tilde{\Gamma}_t\}$ occurs with probability at least $1 - \delta/t^2$. By (14), we have

$$\mathbb{P}\left(a_t \notin \tilde{\Gamma}_t \mid \mathcal{H}_t\right) = \mathbb{P}\left(\mathcal{T} > M_t \mid \mathcal{H}_t\right) = \mathbb{P}\left(\bigcap_{m=1}^{M_t} \left\{a_t^{(m)} \notin \tilde{\Gamma}_t\right\} \mid \mathcal{H}_t\right) = \left(1 - \sum_{i \in \tilde{\Gamma}_t} \tilde{\pi}_i(t)\right)^{M_t}.$$

Since $\gamma < 1/N$, there exists at least one arm in $\tilde{\Gamma}_t$, and thus $\mathbb{P}\left(a_t \notin \tilde{\Gamma}_t \mid \mathcal{H}_t\right) \leq (1 - \gamma)^{M_t}$. If we set M_t as a minimum integer that exceeds $\left(\log \frac{t^2}{\delta}\right) \left(\log \frac{1}{1-\gamma}\right)^{-1}$ then (17) holds. Thus, by choosing M_t for each round that satisfies (17), the algorithm finds an arm j such that $\tilde{\pi}_j(t) > \gamma$ in all rounds with high probability.

Selecting an arm from the set $\{i : \tilde{\pi}_i(t) > \gamma\}$ with high probability is crucial in achieving the regret bound of order $\tilde{O}(\phi^{-2}\sqrt{T})$ for two reasons. First, it guarantees that the arm is super-unsaturated and our novel regret decomposition (6) holds to achieve our regret bound. Let N_t be the set of super-unsaturated arm defined in (5). With Lemma 2, we prove that if $\tilde{\pi}_i(t) > \gamma$ then $i \in N_t$, which implies $\tilde{\Gamma}_t \subseteq N_t$, and thus

$$\mathbb{P}(a_t \in N_t \mid \mathcal{H}_t) \geq \mathbb{P}\left(a_t \in \tilde{\Gamma}_t \mid \mathcal{H}_t\right).$$

Thus we can conclude that a_t is super-unsaturated with probability at least $1 - \delta/t^2$ with M_t defined in Section A.3. Second, the inverse probability, $\pi_{a_t}(t)^{-1}$ is bounded by γ^{-1} which appears in $Y_i^{DR}(t)$ and the proof of Theorem 3. From (15) we can deduce $\pi_{a_t}(t) \geq \tilde{\pi}_{a_t}(t) > \gamma$, for $a_t \in \tilde{\Gamma}_t$. This shows that the assumptions regarding $\pi_{a_t}(t)$ in Theorem 3 hold.

B Technical Lemmas

Lemma 8. [Wainwright, 2019, Theorem 2.19] (Bernstein Concentration) Let $\{D_k, \mathfrak{S}_k\}_{k=1}^\infty$ be a martingale difference sequence and suppose D_k is σ -sub-Gaussian in an adapted sense, i.e. for all $\lambda \in \mathbb{R}$, $\mathbb{E}\left[e^{\lambda D_k} \mid \mathfrak{S}_{k-1}\right] \leq e^{\lambda^2 \sigma^2 / 2}$ almost surely. Then for all $x \geq 0$,

$$\mathbb{P}\left(\left|\sum_{k=1}^n D_k\right| \geq x\right) \leq 2 \exp\left(-\frac{x^2}{2n\sigma^2}\right).$$

Lemma 9. [Azuma, 1967] (Azuma-Hoeffding inequality) If a super-martingale $(Y_t; t \geq 0)$ corresponding to filtration \mathcal{F}_t , satisfies $|Y_t - Y_{t-1}| \leq c_t$ for some constant c_t , for all $t = 1, \dots, T$, then for any $a \geq 0$,

$$\mathbb{P}(Y_T - Y_0 \geq a) \leq e^{-\frac{a^2}{2 \sum_{t=1}^T c_t^2}}.$$

Lemma 10. [Lee et al., 2016, Lemma 2.3] Let $\{N_t\}$ be a martingale on a Hilbert space $(\mathcal{H}, \|\cdot\|_{\mathcal{H}})$. Then there exists a \mathbb{R}^2 -valued martingale $\{P_t\}$ such that for any time $t \geq 0$, $\|P_t\|_2 = \|N_t\|_{\mathcal{H}}$ and $\|P_{t+1} - P_t\|_2 = \|N_{t+1} - N_t\|_{\mathcal{H}}$.

Lemma 11. [Chung and Lu, 2006, Lemma 1, Theorem 32] For a filtration $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \dots \subset \mathcal{F}_T$, suppose each random variable X_t is \mathcal{F}_t -measurable martingale, for $0 \leq t \leq T$. Let B_t denote the bad set associated with the following admissible condition:

$$|X_t - X_{t-1}| \leq c_t,$$

for $1 \leq t \leq T$, where c_1, \dots, c_n are non-negative numbers. Then there exists a collection of random variables Y_0, \dots, Y_T such that Y_t is \mathcal{F}_t -measurable martingale such that

$$|Y_t - Y_{t-1}| \leq c_t,$$

and $\{\omega : Y_t(\omega) \neq X_t(\omega)\} \subset B_t$, for $0 \leq t \leq T$.

Lemma 12. Suppose a random variable X satisfies $\mathbb{E}[X] = 0$, and let η be an R -sub-Gaussian random variable. If $|X| \leq |\eta|$ almost surely, then X is CR -sub-Gaussian for some absolute constant $C > 0$.

Proof. By Proposition 2.5.2 in Vershynin [2018], there exists an absolute constant $C_1 > 0$ such that

$$\mathbb{E} \exp(\lambda^2 \eta^2) \leq \exp\left(\frac{\lambda^2 C_1^2 \sigma^2}{2}\right), \quad \forall \lambda \in \left[-\frac{\sqrt{2}}{C_1 R}, \frac{\sqrt{2}}{C_1 R}\right].$$

Since $|X| \leq |\eta|$ almost surely,

$$\mathbb{E} \exp(\lambda^2 X^2) \leq \exp\left(\frac{\lambda^2 C_1^2 \sigma^2}{2}\right), \quad \forall \lambda \in \left[-\frac{\sqrt{2}}{C_1 R}, \frac{\sqrt{2}}{C_1 R}\right].$$

Since $\mathbb{E}[X] = 0$, by Proposition 2.5.2 in Vershynin [2018], there exists an absolute constant $C_2 > 0$ such that

$$\mathbb{E} \exp(\lambda X) \leq \exp\left(\frac{\lambda^2 C_1^2 C_2^2 \sigma^2}{2}\right), \quad \forall \lambda \in \mathbb{R}.$$

Setting $C = C_1 C_2$ completes the proof. \square

C Missing details in proof of Theorem 1

In section A.3, we prove that $a_t \in \tilde{\Gamma}_t$ with probability at least $1 - \delta/t^2$, for all $t \geq 2$. Thus, for any $x > 0$,

$$\begin{aligned} \mathbb{P}(R(T) > x) &\leq \mathbb{P}\left(R(T) > x, \bigcap_{t=2}^T \{a_t \in \tilde{\Gamma}_t\}\right) + \mathbb{P}\left(\bigcup_{t=2}^T \{a_t \notin \tilde{\Gamma}_t\}\right) \\ &\leq \mathbb{P}\left(R(T) > x, \bigcap_{t=2}^T \{a_t \in \tilde{\Gamma}_t\}\right) + \delta \\ &\leq \mathbb{P}\left(2 + \sum_{t=2}^T \text{regret}(t) > x, \bigcap_{t=2}^T \{a_t \in \tilde{\Gamma}_t\}\right) + \delta \end{aligned}$$

The last inequality holds by Assumption 1. Since $\tilde{\Gamma}_t$ is a subset of N_t and by (6),

$$\begin{aligned} &\mathbb{P}(R(T) > x) \\ &\leq \mathbb{P}\left(2 + \sum_{t=2}^T \left\{2 \|\hat{\beta}_{t-1} - \beta\|_2 + \sqrt{\|X_{a_t^*}(t)\|_{V_{t-1}}^2 + \|X_{a_t}(t)\|_{V_{t-1}}^2}\right\} > x, \bigcap_{t=2}^T \{a_t \in \tilde{\Gamma}_t\}\right) + \delta. \end{aligned} \tag{18}$$

To bound the term $\|\hat{\beta}_t - \beta\|_2$ for all $t = 1, \dots, T-1$, we use Theorem 3. Before that, we need to verify whether the two assumptions on $\pi_i(t)$ in Theorem 3 hold.

First, we show that $\pi_{a_t}(t) > \gamma$. When $t = 1$, we have $\tilde{\pi}_i(1) = 1/N$ for all i . Since $\gamma < 1/N$, we do not need resampling and thus $\pi_i(t) = \tilde{\pi}_i(t) > \gamma$. When $t \geq 2$, $a_t \in \tilde{\Gamma}_t$ is already concerned in (18), and thus $\tilde{\pi}_{a_t}(t) > \gamma$. From (15), we can deduce that $\pi_i(t) > \tilde{\pi}_i(t)$ for all $i \in \tilde{\Gamma}_t$, and thus $\pi_{a_t}(t) > \gamma$.

Now, we prove that $\pi_i(t) > 0$ for all i and t . The case of $t = 1$ is already proved above. When $t \geq 2$, from (15), we have

$$\pi_i(t) := \mathbb{P}(a_t = i | \mathcal{H}_t) = \tilde{\pi}_i(t) \sum_{m=1}^{M_t} \left(1 - \sum_{i \in \tilde{\Gamma}_t} \tilde{\pi}_i(t) \right)^{m-1} > \tilde{\pi}_i(t) > \gamma,$$

for all $i \in \tilde{\Gamma}_t$. If there exists an arm $i \notin \tilde{\Gamma}_t$, from (16),

$$\pi_i(t) = \left(1 - \sum_{i \in \tilde{\Gamma}_t} \tilde{\pi}_i(t) \right)^{M_t-1} \tilde{\pi}_i(t).$$

The first term is positive since there exists an arm $i \notin \tilde{\Gamma}_t$. The second term is also positive since the distribution of $\tilde{\beta}_i(t)$ has support \mathbb{R}^d , which implies that

$$\tilde{\pi}_i(t) := \mathbb{P} \left(X_i(t)^T \tilde{\beta}_i(t) = \max_j X_j(t)^T \tilde{\beta}_j(t) \middle| \mathcal{H}_t \right) > 0,$$

for all i . Thus, $\pi_i(t) > 0$ for all i and t . This implies that the two assumptions on $\pi_i(t)$ in Theorem 3 hold.

Now we can use Theorem 3 and Lemma 6 to have

$$\left\| \hat{\beta}_{t-1} - \beta \right\|_2 \leq \frac{C_{b,\sigma}}{\phi^2 \sqrt{t-1}} \sqrt{\log \frac{12(t-1)^2}{\delta}}, \quad \sqrt{\|X_{a_t^*}(t)\|_{V_{t-1}^{-1}}^2 + \|X_{a_t}(t)\|_{V_{t-1}^{-1}}^2} \leq \frac{1}{\phi \sqrt{N(t-1)}},$$

for all $t = 2, \dots, T$ with probability at least $1 - \delta$. Thus, setting

$$x = 2 + \frac{4C_{b,\sigma}}{\phi^2} \sqrt{T \log \frac{12T^2}{\delta}} + \frac{2\sqrt{T}}{\phi \sqrt{N}}$$

in (18) proves the result.

D Proof of Lemma 2

Proof. First, we bring attention to the fact that the optimal arm a_t^* is in N_t by definition. Suppose that the estimated reward of the optimal arm, $\tilde{Y}_{a_t^*}(t)$ is greater than $\tilde{Y}_j(t)$ for all $j \notin N_t$. In this case, any arm $j \notin N_t$ cannot be the $m_t := \arg \max_i \tilde{Y}_i(t)$. Then we have

$$\begin{aligned} \mathbb{P}(m_t \in N_t | \mathcal{H}_t) &\geq \mathbb{P} \left(\tilde{Y}_{a_t^*}(t) > \tilde{Y}_j(t), \forall j \notin N_t \middle| \mathcal{H}_t \right) \\ &= \mathbb{P} \left(Z_j(t) > \{X_j(t) - X_{a_t^*}(t)\}^T \hat{\beta}_{t-1}, \forall j \notin N_t \middle| \mathcal{H}_t \right), \end{aligned}$$

where $Z_j(t) := \tilde{Y}_{a_t^*}(t) - \tilde{Y}_j(t) - \{X_{a_t^*}(t) - X_j(t)\}^T \hat{\beta}_{t-1}$. Note that $Z_j(t)$ is a Gaussian random variable with mean 0 and variance $v^2(\|X_{a_t^*}(t)\|_{V_{t-1}^{-1}}^2 + \|X_j(t)\|_{V_{t-1}^{-1}}^2)$ given \mathcal{H}_t . For all $j \notin N_t$,

$$\begin{aligned} \{X_j(t) - X_{a_t^*}(t)\}^T \hat{\beta}_{t-1} &= \{X_j(t) - X_{a_t^*}(t)\}^T \{\hat{\beta}_{t-1} - \beta\} - \Delta_j(t) \\ &\leq 2 \left\| \hat{\beta}_{t-1} - \beta \right\|_2 - \Delta_j(t) \leq -\sqrt{\|X_{a_t^*}(t)\|_{V_{t-1}^{-1}}^2 + \|X_j(t)\|_{V_{t-1}^{-1}}^2}. \end{aligned}$$

The last inequality is due to $j \notin N_t$. Thus, we can conclude that

$$\begin{aligned} \mathbb{P}(m_t \in N_t | \mathcal{H}_t) &\geq \mathbb{P} \left(\frac{Z_j(t)}{v \sqrt{\|X_{a_t^*}(t)\|_{V_{t-1}^{-1}}^2 + \|X_j(t)\|_{V_{t-1}^{-1}}^2}} > -\frac{1}{v}, \forall j \notin N_t \middle| \mathcal{H}_t \right) \\ &:= \mathbb{P}(Y_j > -v^{-1}, \forall j \notin N_t | \mathcal{H}_t). \end{aligned}$$

Using the fact that

$$Y_j := \frac{Z_j(t)}{v \sqrt{\|X_{a_t^*}(t)\|_{V_t^{-1}}^2 + \|X_j(t)\|_{V_t^{-1}}^2}}$$

is a standard Gaussian random variable given \mathcal{H}_t , we have

$$\mathbb{P}(Y_j \leq -v^{-1} | \mathcal{H}_t) \leq \exp\left(-\frac{1}{2v^2}\right).$$

Setting $v = \{2 \log(N/(1 - \gamma N))\}^{-1/2}$ gives

$$\mathbb{P}(Y_j \leq -v^{-1} | \mathcal{H}_t) = \exp(-\log(N/(1 - \gamma N))) = \frac{1 - \gamma N}{N}.$$

Thus,

$$\begin{aligned} \mathbb{P}(m_t \in N_t | \mathcal{H}_t) &\geq 1 - \mathbb{P}(Y_j \leq -v^{-1}, \exists j \neq N_t | \mathcal{H}_t) \\ &\geq 1 - \sum_{j \neq N_t} \mathbb{P}(Y_j < -v^{-1} | \mathcal{H}_t) \\ &\geq 1 - (1 - \gamma N) \\ &= \gamma N \\ &\geq 1 - \gamma. \end{aligned}$$

The last inequality holds due to $\gamma \geq 1/(N + 1)$. \square

E Proof of Theorem 3

Proof. Fix $t = 1, \dots, T$ and let $V_t := \sum_{\tau=1}^t \sum_{i=1}^N X_i(\tau) X_i(\tau)^T + \lambda_t I$. For each i and τ , let $\hat{\eta}_i(\tau) = Y_i^{DR}(\tau) - X_i(\tau)^T \beta$. Then

$$\hat{\beta}_t = \beta + V_t^{-1} \left(-\lambda_t \beta + \sum_{\tau=1}^t \sum_{i=1}^N \hat{\eta}_i(\tau) X_i(\tau) \right).$$

To bound $\|\hat{\beta}_t - \beta\|_2$,

$$\begin{aligned} \|\hat{\beta}_t - \beta\|_2 &= \left\| V_t^{-1} \left(-\lambda_t \beta + \sum_{\tau=1}^t \sum_{i=1}^N \hat{\eta}_i(\tau) X_i(\tau) \right) \right\|_2 \\ &\leq \|V_t^{-1}\|_2 \left\| \left(-\lambda_t \beta + \sum_{\tau=1}^t \sum_{i=1}^N \hat{\eta}_i(\tau) X_i(\tau) \right) \right\|_2 \\ &= \{\lambda_{\min}(V_t)\}^{-1} \left\| \left(-\lambda_t \beta + \sum_{\tau=1}^t \sum_{i=1}^N \hat{\eta}_i(\tau) X_i(\tau) \right) \right\|_2. \end{aligned}$$

By Assumption 1, $\|\beta\|_2 \leq 1$. Using triangle inequality,

$$\|\hat{\beta}_t - \beta\|_2 \leq \{\lambda_{\min}(V_t)\}^{-1} \lambda_t + \{\lambda_{\min}(V_t)\}^{-1} \left\| \sum_{\tau=1}^t \sum_{i=1}^N \hat{\eta}_i(\tau) X_i(\tau) \right\|_2. \quad (19)$$

We will bound the first term in (19). Let $\text{Tr}(A)$ be the trace of a matrix A . By the definition of the Frobenious norm, for $\tau = 1, \dots, t$, and for $i = 1, \dots, N$,

$$\left\| \sum_{i=1}^N X_i(\tau) X_i(\tau)^T \right\|_F \leq \sum_{i=1}^N \sqrt{\text{Tr}(X_i(\tau) X_i(\tau)^T X_i(\tau) X_i(\tau)^T)} \leq N.$$

By Assumptions 3 and 4, $\left\{ \sum_{i=1}^N X_i(\tau) X_i(\tau)^T \right\}_{\tau=1}^t$ are independent random variables such that $\mathbb{E} \left[\sum_{i=1}^N X_i(\tau) X_i(\tau)^T \right] \geq N\phi^2 > 0$. Let $\delta \in (0, 1)$ be given. By Lemma 6, if we set $\lambda_t = 4\sqrt{2}N\sqrt{t \log \frac{12t^2}{\delta}}$,

$$\{\lambda_{\min}(V_i)\}^{-1} < \frac{1}{\phi^2 N t},$$

holds with probability at least $1 - \delta/(3t^2)$. Thus, the first term can be bounded by

$$\{\lambda_{\min}(V_i)\}^{-1} \lambda_t \leq \frac{4\sqrt{\log \frac{12t^2}{\delta}}}{\sqrt{t}\phi^2}. \quad (20)$$

Now we will bound the second term in (19). Let $U_i(\tau) := X_i(\tau) X_i(\tau)^T (\check{\beta}_\tau - \beta)$. Then we can decompose $\hat{\eta}_i(\tau) X_i(\tau)$ as,

$$\begin{aligned} \hat{\eta}_i(\tau) X_i(\tau) &= U_i(\tau) + \frac{\mathbb{I}(a_\tau = i)}{\pi_i(\tau)} \left(Y_i(\tau) - X_i(\tau)^T \check{\beta}_\tau \right) X_i(\tau) \\ &= \left(1 - \frac{\mathbb{I}(a_\tau = i)}{\pi_i(\tau)} \right) U_i(\tau) + \frac{\mathbb{I}(a_\tau = i)}{\pi_i(\tau)} \eta_i(\tau) X_i(\tau) \\ &:= D_i(\tau) + E_i(\tau). \end{aligned} \quad (21)$$

Let $D_\tau := \sum_{i=1}^N D_i(\tau)$. Since $U_i(\tau)$ is \mathcal{H}_τ -measurable, the conditional expectation of D_τ is

$$\begin{aligned} \mathbb{E}[D_\tau | \mathcal{H}_\tau] &= \mathbb{E} \left[\sum_{i=1}^N D_i(\tau) \middle| \mathcal{H}_\tau \right] = \sum_{i=1}^N \mathbb{E} \left[\left(1 - \frac{\mathbb{I}(a_\tau = i)}{\pi_i(\tau)} \right) \middle| \mathcal{H}_\tau \right] U_i(\tau) \\ &= \sum_{i=1}^N \left(1 - \frac{\pi_i(\tau)}{\pi_i(\tau)} \right) U_i(\tau) = 0 \end{aligned}$$

Thus, $\{ \sum_{u=1}^\tau D_u \}_{\tau=1}^t$ is a martingale sequence on $(\mathbb{R}^d, \|\cdot\|_2)$ with respect to \mathcal{H}_τ . By Lemma 10, since $(\mathbb{R}^d, \|\cdot\|_2)$ is a Hilbert space, there exists a martingale sequence $\{P_\tau\}_{\tau=1}^t = \left\{ (P_\tau^{(1)}, P_\tau^{(2)})^T \right\}_{\tau=1}^t$ on \mathbb{R}^2 such that

$$\left\| \sum_{u=1}^\tau D_u \right\|_2 = \|P_\tau\|_2, \quad \|D_\tau\|_2 = \|P_\tau - P_{\tau-1}\|_2 \quad (22)$$

and $P_0 = 0$, for any $\tau = 1, \dots, t$. Since $\|\check{\beta}_\tau - \beta\|_2 \leq b$, for $r = 1, 2$

$$\begin{aligned} \left| P_\tau^{(r)} - P_{\tau-1}^{(r)} \right| &\leq \|P_\tau - P_{\tau-1}\|_2 = \|D_\tau\|_2 \\ &= \left\| \sum_{i=1}^N \left(1 - \frac{\mathbb{I}(a_\tau = i)}{\pi_i(\tau)} \right) U_i(\tau) \right\|_2 \\ &\leq \sum_{i=1}^N \left| 1 - \frac{\mathbb{I}(a_\tau = i)}{\pi_i(\tau)} \right| \|U_i(\tau)\|_2 \\ &\leq \sum_{i=1}^N \left| 1 - \frac{\mathbb{I}(a_\tau = i)}{\pi_i(\tau)} \right| \|\check{\beta}_\tau - \beta\|_2 \\ &\leq \left(N - 1 + \frac{1}{\pi_{a_\tau}(\tau)} - 1 \right) b \\ &\leq (N + \pi_{a_\tau}(\tau)^{-1}) b. \end{aligned}$$

By Lemma 11, there exists a martingale sequence $\{N_\tau^{(r)}\}_{\tau=1}^t$ such that $|N_\tau^{(r)} - N_{\tau-1}^{(r)}| \leq (N + \gamma^{-1})b$, for all $\tau = 1, \dots, t$ and

$$\{N_t^{(r)} \neq P_t^{(r)}\} \subset \bigcup_{\tau=1}^t \{|P_\tau^{(r)} - P_{\tau-1}^{(r)}| > (N + \gamma^{-1})b\} \subset \bigcup_{\tau=1}^t \{\pi_{a_\tau}(\tau) \leq \gamma\}. \quad (23)$$

Thus, by (22) and (23), for any $x > 0$,

$$\begin{aligned} \mathbb{P}\left(\left\|\sum_{u=1}^t D_u\right\|_2 > x, \bigcap_{\tau=1}^T \{\pi_{a_\tau}(\tau) > \gamma\}\right) &= \mathbb{P}\left(\|P_t\|_2 \geq x, \bigcap_{\tau=1}^T \{\pi_{a_\tau}(\tau) > \gamma\}\right) \\ &\leq \mathbb{P}\left(\sum_{r=1}^2 |P_t^{(r)}| \geq x, \bigcap_{\tau=1}^t \{\pi_{a_\tau}(\tau) > \gamma\}\right) \\ &\leq \sum_{r=1}^2 \mathbb{P}\left(|P_t^{(r)}| \geq \frac{x}{2}, \bigcap_{\tau=1}^t \{\pi_{a_\tau}(\tau) > \gamma\}\right) \\ &\leq \sum_{r=1}^2 \mathbb{P}\left(|P_t^{(r)}| \geq \frac{x}{2}, N_t^{(r)} = P_t^{(r)}\right) \\ &\leq \sum_{r=1}^2 \mathbb{P}\left(|N_t^{(r)}| \geq \frac{x}{2}\right). \end{aligned}$$

Since $N_\tau^{(r)}$ has bounded differences, we can apply Lemma 9 to have

$$\sum_{r=1}^2 \mathbb{P}\left(|N_t^{(r)}| \geq \frac{x}{2}\right) \leq 4 \exp\left(-\frac{x^2}{8tb^2(N + \gamma^{-1})^2}\right).$$

Thus, with probability at least $1 - \delta/(3t^2)$,

$$\left\|\sum_{\tau=1}^t D_\tau\right\|_2 \leq 2\sqrt{2}b(N + \gamma^{-1})\sqrt{\log \frac{12t^2}{\delta}} \quad (24)$$

holds with the event $\bigcap_{t=1}^T \{\pi_{a_t}(t) > \gamma\}$.

Now we will bound the $E_i(\tau)$ term in (21). Under the event $\bigcap_{t=1}^T \{\pi_{a_t}(t) > \gamma\}$, we have

$$\sum_{\tau=1}^t \sum_{i=1}^N E_i(\tau) = \sum_{\tau=1}^t \frac{\eta_{a_\tau}(\tau)}{\pi_{a_\tau}(\tau)} X_{a_\tau}(\tau) = \sum_{\tau=1}^t \frac{\mathbb{I}(\pi_{a_t}(t) > \gamma) \eta_{a_\tau}(\tau)}{\pi_{a_\tau}(\tau)} X_{a_\tau}(\tau)$$

For each $\tau \geq 1$, define a filtration $\mathcal{F}_{\tau-1} := \mathcal{H}_\tau \cup \{a_\tau\}$. Then $X_{a_\tau}(\tau)$ is $\mathcal{F}_{\tau-1}$ -measurable. By Assumption 2, for any $\lambda \in \mathbb{R}$,

$$\mathbb{E}\left[\exp\left(\lambda \frac{\mathbb{I}(\pi_{a_t}(t) > \gamma) \eta_{a_\tau}(\tau)}{\pi_{a_\tau}(\tau)}\right) \middle| \mathcal{F}_{\tau-1}\right] \leq \exp\left(\frac{\lambda^2 \mathbb{I}(\pi_{a_t}(t) > \gamma) \sigma^2}{2\pi_{a_\tau}(\tau)^2}\right) \leq \exp\left(\frac{\lambda^2 \sigma^2}{2\gamma^2}\right),$$

almost surely. Since $\|X_{a_\tau}(\tau)\|_2 \leq 1$, by Lemma 4, there exists an absolute constant $C > 0$ such that, with probability at least $1 - \delta/(3t^2)$,

$$\left\|\sum_{\tau=1}^t \sum_{i=1}^N E_i(\tau)\right\|_2 \leq 2C\sigma\gamma^{-1}\sqrt{t}\sqrt{\log \frac{12t^2}{\delta}}. \quad (25)$$

Thus, with (20), (24), and (25), under the event $\bigcap_{t=1}^T \{\pi_{a_t}(t) > \gamma\}$, we have

$$\begin{aligned}
\|\widehat{\beta}_t - \beta\|_2 &\leq \frac{4\sqrt{\log \frac{12t^2}{\delta}}}{\sqrt{t}\phi^2} + \frac{1}{\phi^2 N t} \left(4(N + \gamma^{-1}) b \sqrt{t} \sqrt{\log \frac{12t^2}{\delta}} + 2C\sigma\gamma^{-1} \sqrt{t} \sqrt{\log \frac{12t^2}{\delta}} \right) \\
&\leq \frac{4 + 4b + \gamma^{-1} N^{-1} (4b + 2C\sigma)}{\phi^2 \sqrt{t}} \sqrt{\log \frac{12t^2}{\delta}} \\
&\leq \frac{4 + 4b + 2(4b + 2C\sigma)}{\phi^2 \sqrt{t}} \sqrt{\log \frac{12t^2}{\delta}} \\
&:= \frac{C_{b,\sigma}}{\phi^2 \sqrt{t}} \sqrt{\log \frac{12t^2}{\delta}},
\end{aligned} \tag{26}$$

with probability at least $1 - \delta/t^2$. Since (26) holds for all $t = 1, \dots, T$,

$$\begin{aligned}
&\leq \mathbb{P} \left(\bigcup_{t=1}^T \left\{ \|\widehat{\beta}_t - \beta\|_2 > \frac{C_{b,\sigma}}{\phi^2 \sqrt{t}} \sqrt{\log \frac{12t^2}{\delta}} \right\}, \bigcap_{t=1}^T \{\pi_{a_t}(t) > \gamma\} \right) \\
&\leq \mathbb{P} \left(\bigcup_{t=1}^T \left\{ \|\widehat{\beta}_t - \beta\|_2 > \frac{C_{b,\sigma}}{\phi^2 \sqrt{t}} \sqrt{\log \frac{12t^2}{\delta}} \right\}, \bigcap_{t=1}^T \{\pi_{a_t}(t) > \gamma\} \right) \\
&\leq \sum_{t=1}^T \mathbb{P} \left(\|\widehat{\beta}_t - \beta\|_2 > \frac{C_{b,\sigma}}{\phi^2 \sqrt{t}} \sqrt{\log \frac{12t^2}{\delta}}, \bigcap_{t=1}^T \{\pi_{a_t}(t) > \gamma\} \right) \\
&\leq \delta.
\end{aligned}$$

□

F Proof of Lemma 4

Proof. Fix a $t \geq 1$. Since for each $\tau = 1, \dots, t$, $\mathbb{E}[\eta(\tau) | \mathcal{F}_{\tau-1}] = 0$ and $X(\tau)$ is $\mathcal{F}_{\tau-1}$ -measurable, the stochastic process,

$$\left\{ \sum_{\tau=1}^u \eta(\tau) X(\tau) \right\}_{u=1}^t \tag{27}$$

is a \mathbb{R}^d -martingale. Since $(\mathbb{R}^d, \|\cdot\|_2)$ is a Hilbert space, by Lemma 10, there exists a \mathbb{R}^2 -martingale $\{M_u\}_{u=1}^t$ such that

$$\left\| \sum_{\tau=1}^u \eta(\tau) X(\tau) \right\|_2 = \|M_u\|_2, \quad \|\eta(u) X(u)\|_2 = \|M_u - M_{u-1}\|_2, \tag{28}$$

and $M_0 = 0$. Set $M_u = (M_1(u), M_2(u))^T$. Then for each $i = 1, 2$, and $u \geq 2$, by the assumption $\|X(u)\|_2 \leq 1$,

$$\begin{aligned}
|M_i(u) - M_i(u-1)| &\leq \|M_u - M_{u-1}\|_2 \\
&= \|\eta(u) X(u)\|_2 \\
&\leq |\eta(u)|.
\end{aligned}$$

By Lemma 12, $M_i(u) - M_i(u-1)$ is $C\sigma$ -sub-Gaussian for some constant $C > 0$. By Lemma 9, for $x > 0$,

$$\begin{aligned}
\mathbb{P}(|M_i(t)| > x) &= \mathbb{P} \left(\left| \sum_{u=1}^t M_i(u) - M_i(u-1) \right| > x \right) \\
&\leq 2 \exp \left(-\frac{x^2}{2tC^2\sigma^2} \right),
\end{aligned}$$

for each $i = 1, 2$. Thus, with probability $1 - \delta/(2t^2)$,

$$M_i(t)^2 \leq 2tC^2\sigma^2 \log \frac{4t^2}{\delta}.$$

In summary, with probability at least $1 - \delta/t^2$,

$$\left\| \sum_{\tau=1}^t \eta(\tau)X(\tau) \right\|_2 = \sqrt{M_1(t)^2 + M_2(t)^2} \leq 2C\sigma\sqrt{t} \sqrt{\log \frac{4t^2}{\delta}}.$$

□

G Proof of Lemma 6

Proof. For each $\tau = 1, \dots, t$, let $\Sigma_\tau = \mathbb{E}[P(\tau) | \mathcal{F}_{\tau-1}]$. Since $P(\tau)$ and Σ_τ are symmetric matrices,

$$\begin{aligned} \lambda_{\min} \left(\sum_{\tau=1}^t P(\tau) + \lambda_t I \right) &= \lambda_{\min} \left(\sum_{\tau=1}^t P(\tau) \right) + \lambda_t \\ &= \lambda_{\min} \left(\sum_{\tau=1}^t \{P(\tau) - \Sigma_\tau\} + \sum_{\tau=1}^t \Sigma_\tau \right) + \lambda_t \\ &\geq \lambda_{\min} \left(\sum_{\tau=1}^t \{P(\tau) - \Sigma_\tau\} \right) + \sum_{\tau=1}^t \lambda_{\min}(\Sigma_\tau) + \lambda_t \\ &\geq \lambda_{\min} \left(\sum_{\tau=1}^t \{P(\tau) - \Sigma_\tau\} \right) + \phi^2 t + \lambda_t. \end{aligned}$$

The last inequality uses the fact that $\lambda_{\min}(\Sigma_\tau) \geq \phi^2$ for all τ .

$$\begin{aligned} \mathbb{P} \left(\lambda_{\min} \left(\sum_{\tau=1}^t P(\tau) + \lambda_t I \right) \leq \phi^2 t \right) &\leq \mathbb{P} \left(\lambda_{\min} \left(\sum_{\tau=1}^t \{P(\tau) - \Sigma_\tau\} \right) + \lambda_t \leq 0 \right) \\ &= \mathbb{P} \left(\lambda_{\max} \left(\sum_{\tau=1}^t \{\Sigma_\tau - P(\tau)\} \right) \geq \lambda_t \right) \\ &\leq \mathbb{P} \left(\left\| \sum_{\tau=1}^t \{\Sigma_\tau - P(\tau)\} \right\|_F \geq \lambda_t \right). \end{aligned} \quad (29)$$

Set $S_u = \sum_{\tau=1}^u \{\Sigma_\tau - P(\tau)\}$. Then $\{S_u\}_{u=1}^t$ can be regarded as a martingale sequence on $\mathbb{R}^{d \times d}$ with respect to $\{P(\tau)\}_{\tau=1}^t$. Note that $(\mathbb{R}^{d \times d}, \|\cdot\|_F)$ is a Hilbert space. By Lemma 10, there exists a martingale sequence $\{D_u = (D_1(u), D_2(u))^T\}_{u=1}^t$ on \mathbb{R}^2 such that

$$\|S_u\|_F = \sqrt{D_1(u)^2 + D_2(u)^2}, \quad \|M_u - \Sigma_u\|_F = \|D_u - D_{u-1}\|_2, \quad (30)$$

for any $u \geq 1$, and $D_0 = 0$. Then, for any $i = 1, 2$,

$$|D_i(u) - D_i(u-1)|^2 \leq \|D_u - D_{u-1}\|_2^2 = \|P(u) - \Sigma_u\|_F^2$$

Since $\|P(u) - \Sigma_u\|_F \leq 2c$, we can apply Lemma 9 for $D_1(\tau)$, and $D_2(\tau)$, respectively. For any $i = 1, 2$, and for any $x > 0$,

$$\mathbb{P}(|D_i(t)| \geq x) \leq 2 \exp \left(-\frac{x^2}{8c^2 t} \right).$$

From (29) and (30),

$$\begin{aligned}
\mathbb{P} \left(\lambda_{\min} \left(\sum_{\tau=1}^t P(\tau) + \lambda_t I \right) \leq \phi^2 t \right) &\leq \mathbb{P} (\|S_t\|_F \geq \lambda_t) \\
&= \mathbb{P} \left(\sqrt{D_1(t)^2 + D_2(t)^2} \geq \lambda_t \right) \\
&\leq \mathbb{P} (|D_1(t)| + |D_2(t)| \geq \lambda_t) \\
&\leq \mathbb{P} \left(|D_1(t)| \geq \frac{\lambda_t}{2} \right) + \mathbb{P} \left(|D_2(t)| \geq \frac{\lambda_t}{2} \right) \\
&\leq 4 \exp \left(-\frac{\lambda_t^2}{32c^2 t} \right).
\end{aligned}$$

Thus, for any $\delta \in (0, 1)$, if $\lambda_t \geq 4\sqrt{2}c\sqrt{t}\sqrt{\log \frac{4t^2}{\delta}}$, then with probability at least $1 - \delta$,

$$\lambda_{\min} \left(\sum_{\tau=1}^t P(\tau) + \lambda_t I \right) > \phi^2 t.$$

□

H Implementation details

H.1 Efficient calculation of the sampling probability

In our proposed algorithm, we use quasi-Monte Carlo estimation to calculate the sampling probability, $\tilde{\pi}_i(t)$. At round t , for each $i = 1, \dots, N$, define $Z_i = \frac{X_i(t)^T (\hat{\beta}_i(t) - \hat{\beta}_{t-1})}{v \|X_i(t)\|_{V_t^{-1}}}$. Then, Z_1, \dots, Z_N are IID standard Gaussian random variables. For each $i = 1, \dots, N$,

$$\begin{aligned}
\tilde{\pi}_i(t) &= \mathbb{P} \left(X_i(t)^T \tilde{\beta}_i(t) \geq X_j(t)^T \tilde{\beta}_j(t), \forall j \neq i \mid \mathcal{H}_t \right) \\
&= \mathbb{P} \left(\frac{\|X_i(t)\|_{V_t^{-1}}}{\|X_j(t)\|_{V_t^{-1}}} Z_i \geq Z_j + \frac{(X_j(t) - X_i(t))^T \hat{\beta}_{t-1}}{v \|X_j(t)\|_{V_t^{-1}}}, \forall j \neq i \mid \mathcal{H}_t \right)
\end{aligned}$$

let f and F be the density and the distribution function of the standard Gaussian random variables, respectively. Since Z_i , and $\{Z_j\}_{j \neq i}$ are independent, the selection probability can be written as,

$$\tilde{\pi}_i(t) = \int \prod_{j \neq i} F \left(\frac{\|X_i(t)\|_{V_t^{-1}}}{\|X_j(t)\|_{V_t^{-1}}} z + \frac{(X_i(t) - X_j(t))^T \hat{\beta}_{t-1}}{v \|X_j(t)\|_{V_t^{-1}}} \right) f(z) dz.$$

This can be estimated by,

$$\frac{1}{M} \sum_{m=1}^M F \prod_{j \neq i} \left(\frac{\|X_i(t)\|_{V_t^{-1}}}{\|X_j(t)\|_{V_t^{-1}}} Z^{(m)} + \frac{(X_i(t) - X_j(t))^T \hat{\beta}_{t-1}}{v \|X_j(t)\|_{V_t^{-1}}} \right), \quad (31)$$

where $Z^{(m)}$ is the standard Gaussian random variables.

In this way, we can compute $\tilde{\pi}_i(t)$ without sampling $\tilde{\beta}_i(t)$ $M \times N$ times from $N(\hat{\beta}_{t-1}, v_t I)$. The error of the quasi Monte Carlo method is bounded by $O \left(\frac{(\log M)^s}{M} \right)$, where s is the dimension of the domain of function to integrate. If we sample $\tilde{\beta}_i(t)$ $M \times N$ times, it gives $O \left(\frac{(\log M)^{N-1}}{M} \right)$ error. In contrast, using (31) reduces the error to $O \left(\frac{\log M}{M} \right)$.

In our simulation studies, we use `sobol_seq` module in Python 3 to generate the quasi-Monte Carlo samples. The number of samples is $M = 200$ in BLTS and DRTS. We plot the estimator of $\tilde{\pi}_i(t)$ using $m = 1, \dots, 200$ quasi-Monte Carlo samples, and observe that it converges within the small errors.

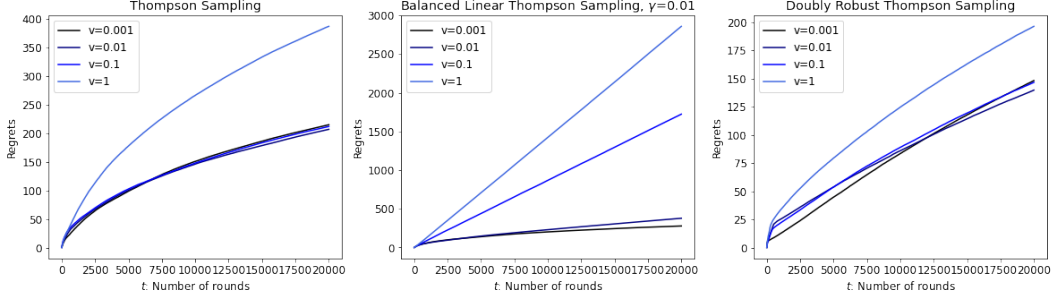


Figure 2: A comparison of the cumulative regrets of LinTS (left), BLTS (middle), and DRTS (right) on various v when $d = 20$ and $N = 20$. Each line shows the averaged cumulative regrets over 10 repeated experiments.

H.2 Simulation results with various hyperparameters

In this subsection, we report the performance of the three algorithms, (i) LinTS, (ii) BLTS, and (iii) the proposed DRTS, with various hyperparameters. As described in Section 5, the hyperparameter sets are $v \in \{0.001, 0.01, 0.1, 1\}$ for all three algorithms and $\gamma \in \{0.01, 0.5, 0.1\}$ for BLTS.

Figure 2 shows the comparison of the three algorithms on various hyperparameter v , when $d = 20$, $N = 20$, $\gamma = 0.01$. We find that the performance of the three algorithms do not change much when $v \leq 0.01$. This trend is similar on different γ , N , and d .

I A review of approaches to missing data and doubly-robust method

In this section, we review approaches to missing data and the doubly-robust method used in our proposed method. First, we provide the approaches from a purely missing data point of view and how the doubly-robust method is motivated. In the second section, we show the procedures applying the doubly-robust method in bandit settings.

I.1 Doubly-robust method in missing data

There are two main approaches to missing data: imputation and inverse probability weighting (IPW). Imputation is to fill in the predicted value of missing data from a specified model, and IPW is to use the observed records only but weight them by the inverse of the observation probability. The doubly-robust method can be viewed as a combination of the two.

For illustrative purposes, consider the problem of estimating the marginal mean of $Y \in \mathbb{R}$, $\mathbb{E}(Y) =: \mu$. Denoting $(Y_i - \mu)$ by $U_i(\mu)$, when all data are observed,

$$U(\mu) = \sum_{i=1}^n U_i(\mu) = \sum_{i=1}^n (Y_i - \mu) = 0,$$

gives an unbiased estimator of μ , $\sum_{i=1}^n Y_i/n$, and $U(\mu)$ is called an unbiased estimating function since $\mathbb{E}[U(\mu)] = 0$. Let δ_i be the observation indicator which takes value 1 if Y_i is observed, 0, otherwise. Suppose there are auxiliary variables, $X_i \in \mathbb{R}^d$, and X_i 's are observed for all i . Also denote the probability of observation by $P(\delta_i = 1|X_i) =: \pi_i$. We assume $P(\delta_i = 1|Y_i, X_i) = P(\delta_i = 1|X_i)$, that is, the observation indicator is independent of Y_i . This is called *missing at random* mechanism. This assumption is required for the doubly robust method to be valid. Using the observed values only, the estimating equation for the observed data

$$U_o(\mu) = \sum_{i=1}^n \delta_i U_i(\mu) = \sum_{i=1}^n \delta_i (Y_i - \mu) = 0,$$

gives $\frac{\sum_{i=1}^n \delta_i Y_i}{\sum_{i=1}^n \delta_i}$ as an estimator for μ . This estimator may be biased since $\mathbb{E}U_o(\mu) \neq 0$.

The two main approaches modify the observed estimating function employing two new quantities, $\mathbb{E}(Y_i|X_i)$ and π_i . These two quantities are usually unknown and we need to specify models. Therefore

the two approaches require assumptions for auxiliary models: the imputation model, $\mathbb{E}(Y_i|X_i; \beta)$, and the model for observation probability, $\pi_i(\phi)$. The validity of each approach depends on the correct specification of the auxiliary model assumptions. The qualifier ‘auxiliary’ comes from the fact that these models are not needed when there is no missing data. In IPW, one constructs an unbiased estimating equation by amplifying the observed record according to the inverse of the observation probability as follows:

$$\sum_{i=1}^n \frac{\delta_i}{\pi_i(\phi)} U_i(\mu) = \sum_{i=1}^n \frac{\delta_i}{\pi_i(\phi)} (Y_i - \mu).$$

If $\pi(\phi)$ is correctly specified, i.e., $\pi = \pi(\phi)$, $\mathbb{E}(\sum_{i=1}^n \frac{\delta_i}{\pi_i(\phi)} U_i(\mu)) = 0$, hence the resulting IPW estimator is valid. In the imputation method, we replace missing Y_i with $\mathbb{E}(Y_i|X_i; \beta)$ and the estimator is the solution of $U^{IMP}(\mu, \beta) = 0$ where

$$\begin{aligned} U^{IMP}(\mu, \beta) &= \sum_{i=1}^n [\delta_i U_i(\mu) + (1 - \delta_i) \mathbb{E}(U_i(\mu)|X_i; \beta)] \\ &= \sum_{i=1}^n [\mathbb{E}(Y_i|X_i; \beta) + \delta_i \{Y_i - \mathbb{E}(Y_i|X_i; \beta)\} - \mu]. \end{aligned}$$

The doubly robust (DR) method [Robins et al., 1994, Bang and Robins, 2005] was initially motivated by attempting to improve the efficiency of the IPW method. Note that we can construct an auxiliary unbiased estimating function $(\frac{\delta_i}{\pi_i(\phi)} - 1)$. Geometrically we can reduce the norm of the estimating function $\frac{\delta_i}{\pi_i(\phi)} U_i(\mu)$ by subtracting the projection on to the nuisance tangent space formed from $(\frac{\delta_i}{\pi_i(\phi)} - 1)$. The nuisance tangent space is the closed linear span of $B(\frac{\delta_i}{\pi_i(\phi)} - 1)$ for some $B \in \mathbb{R}^d$, and the projection onto the nuisance tangent space is

$$\sum_{i=1}^n \frac{\delta_i - \pi_i(\phi)}{\pi_i(\phi)} \mathbb{E}(U_i|X_i; \beta).$$

After subtraction, the DR estimating function has a form

$$\begin{aligned} U^{DR}(\mu, \beta, \phi) &= \sum_{i=1}^n \left[\frac{\delta_i}{\pi_i(\phi)} U_i(\mu) + (1 - \frac{\delta_i}{\pi_i(\phi)}) \mathbb{E}(U_i|X_i; \beta) \right] \\ &= \sum_{i=1}^n \left[\mathbb{E}(U_i|X_i; \beta) + \frac{\delta_i}{\pi_i(\phi)} \{U_i(\mu) - \mathbb{E}(U_i(\mu)|X_i; \beta)\} \right]. \end{aligned}$$

Note that when you replace δ_i in $U^{IMP}(\mu)$ with $\frac{\delta_i}{\pi_i(\phi)}$, you obtain $U^{DR}(\mu)$. The DR method requires both auxiliary models. However, its validity is guaranteed when *either* of the models is correct. To verify, if the imputation model is correctly specified, i.e., $\mathbb{E}[U_i(\mu) - \mathbb{E}(U_i(\mu)|X_i; \beta)|X_i] = 0$, we have

$$\mathbb{E}\{U^{DR}(\mu, \beta, \phi)\} = \mathbb{E} \sum_{i=1}^n \left[\mathbb{E}(U_i|X_i) - \frac{\delta_i}{\pi_i(\phi)} \{U_i(\mu) - \mathbb{E}(U_i(\mu)|X_i)\} \right] = \sum_{i=1}^n \mathbb{E} \mathbb{E}(U_i|X_i) = 0$$

even if the π model is misspecified, i.e., $\pi_i(\phi) \neq \pi_i$. If the observation model is correctly specified, $\pi_i(\phi) = \pi_i$, then $\mathbb{E}(1 - \frac{\delta_i}{\pi_i}|X_i) = 0$, and

$$\mathbb{E}\{U^{DR}(\mu, \beta, \phi)\} = \sum_{i=1}^n \mathbb{E} \left[\frac{\delta_i}{\pi_i} U_i(\mu) + \left\{ (1 - \frac{\delta_i}{\pi_i}) \mathbb{E}(U_i|X_i; \beta) \right\} \right] = \sum_{i=1}^n \mathbb{E} \left[\frac{\delta_i}{\pi_i} U_i(\mu) \right] = 0,$$

even if the imputation model is misspecified, i.e., $\mathbb{E}[U_i(\mu)|X_i] \neq \mathbb{E}[U_i(\mu)|X_i; \beta]$. Therefore when *either* of the models is correct, $U^{DR}(\mu)$ is unbiased and with other technical conditions, the estimator can be shown to be consistent. That is why the qualifier *doubly robust* is adopted. The construction of the DR estimating function is possible because we have two unbiased estimating functions.

I.2 Application to bandit settings

In bandit settings, the missingness is controlled since the learner selects the arm. Therefore, the probability of observation or selection is known and the DR estimator is guaranteed to be valid although the imputation model for missing reward is incorrectly specified. The merit of the DR estimator in the bandit setting is that we can utilize the observed contexts from selected or unselected arms. Below we describe the DR method in the contextual bandit setting.

Let $\pi_i(t) := \mathbb{P}(a_t = i | \mathcal{H}_t)$ be the probability of selecting arm i at round t . As defined in the manuscript, the DR pseudo-reward is

$$Y_i^{DR}(t) = \left\{ 1 - \frac{\mathbb{I}(i = a_t)}{\pi_i(t)} \right\} X_i(t)^T \check{\beta}_t + \frac{\mathbb{I}(i = a_t)}{\pi_i(t)} Y_{a_t}(t), \quad (32)$$

for some $\check{\beta}_t$ depending on \mathcal{H}_t . The pseudo-reward (32) comes from the following procedures. First we construct an unbiased estimating function also known as the IPW score,

$$\sum_{\tau=1}^t \sum_{i=1}^N \frac{\mathbb{I}(i = a_\tau)}{\pi_i(\tau)} X_i(\tau) (Y_i(\tau) - X_i(\tau)^T \beta), \quad (33)$$

where only the pairs $(X_i(t), Y_i(t))$ from the selected arms are contributed according the weight of the inverse of $\pi_i(t)$. Setting this score equal to 0 and solving β gives the estimator used in Dimakopoulou et al. [2019]. Now we can subtract the projection on the nuisance tangent space from (33). The nuisance tangent space is the closed linear span of $B(\frac{\mathbb{I}(i=a(t))}{\pi_i(t)} - 1)$ for some $B \in \mathbb{R}^d$, and the projection onto the nuisance tangent space is

$$\sum_{\tau=1}^t \sum_{i=1}^N \frac{\mathbb{I}(i = a_\tau) - \pi_i(\tau)}{\pi_i(\tau)} X_i(\tau) (E(Y_i(\tau) | \mathcal{H}_\tau) - X_i(\tau)^T \beta).$$

When the projection is subtracted from the (33) after replacing $E(Y_i(t) | \mathcal{H}_t)$ with $X_i(t)^T \check{\beta}_t$, the IPW score becomes the efficient score,

$$\sum_{\tau=1}^t \sum_{i=1}^N X_i(\tau) (Y_i^{DR}(\tau) - X_i(\tau)^T \beta). \quad (34)$$

Any $\check{\beta}_t$ that depends on \mathcal{H}_t serves the purpose of imputation. Due to the doubly robustness property, $X_i(t)^T \check{\beta}_t$ does not have to be an unbiased estimator of $E(Y_i(t) | \mathcal{H}_t)$. We recommend setting $\check{\beta}_t$ as the ridge regression estimator based on the selected arms only. The expression (34) resembles the score when the rewards for all arms were observed, if $Y_i(t)$ is replaced with $Y_i^{DR}(t)$.

Our proposed estimator $\hat{\beta}_t$ is a solution of (34) with a regularization parameter λ_t :

$$\hat{\beta}_t = \left(\sum_{\tau=1}^t \sum_{i=1}^N X_i(\tau) X_i(\tau)^T + \lambda_t I \right)^{-1} \left(\sum_{\tau=1}^t \sum_{i=1}^N X_i(\tau) Y_i^{DR}(\tau) \right).$$

Harnessing the pseudo-rewards defined in (32), we can make use of all contexts rather than just selected contexts. The use of all contexts instead of $X_{a_t}(t)$ induces the improvement in the regret bound of the proposed algorithm. Kim and Paik [2019] also suggests DR estimator, but it uses Lasso estimator from the following pseudo-reward

$$Y_i^{DR}(t) = \bar{X}(t)^T \hat{\beta}(t-1) + \frac{1}{N} \frac{Y_{a(t)}(t) - b_{a(t)}(t)^T \hat{\beta}(t-1)}{\pi_{a(t)}(t)},$$

where $\bar{X}(t) = \frac{1}{N} \sum_{i=1}^N X_i(t)$. This estimator is of an aggregated form. As described in the text, the estimator using the aggregated pseudo-reward does not permit the regret decomposition as equation (6) in the paper.

J Limitations of our work

1. The regret bound is constructed under the assumption that the contexts are independent over rounds (Assumption 3) and the covariance matrix is positive definite (Assumption 4). When using our proposed algorithm, one should check that the contexts satisfies the two assumptions. When the contexts violates the two assumptions, the improved regret bound (3) might not hold.
2. Our proposed algorithm, DRTS requires additional computations for $\tilde{\pi}_i(t)$, $\pi_i(t)$ and the imputation estimator $\check{\beta}$. To lessen this computational burdens we developed an efficient way to compute $\tilde{\pi}_i(t)$ and $\pi_i(t)$.