
A Simple Unified Framework for High Dimensional Bandit Problems

Wenjie Li¹ Adarsh Barik² Jean Honorio²

Abstract

Stochastic high dimensional bandit problems with low dimensional structures are useful in different applications such as online advertising and drug discovery. In this work, we propose a simple unified algorithm for such problems and present a general analysis framework for the regret upper bound of our algorithm. We show that under some mild unified assumptions, our algorithm can be applied to different high-dimensional bandit problems. Our framework utilizes the low dimensional structure to guide the parameter estimation in the problem, therefore our algorithm achieves the comparable regret bounds in the LASSO bandit as a sanity check, as well as novel bounds that depend logarithmically on dimensions in the low-rank matrix bandit, the group sparse matrix bandit, and in a new problem: the multi-agent LASSO bandit.

1. Introduction

Stochastic multiarmed contextual bandits are useful models in various application domains, such as recommendation systems, online advertising, and personalized healthcare (Auer, 2002b; Chu et al., 2011; Abbasi-Yadkori et al., 2011). Under this setting, the agent chooses one specific arm at each round and observes a reward, which is modeled as a function of an unknown parameter and the context of the arm. In practice, such problems are often high-dimensional, but the unknown parameter is typically assumed to have low-dimensional structure, which in turns implies a succinct representation of the final reward.

For example, in high dimensional sparse linear bandits, also known as the LASSO bandit problem (Bastani & Bayati, 2020), both the contexts and the unknown parameter are

high-dimensional vectors, while the parameter is assumed to be sparse with limited nonzero elements. There has been a line of research on LASSO bandits, such as Abbasi-Yadkori et al. (2011), Alexandra Carpentier (2012), Bastani & Bayati (2020), Lattimore et al. (2015), Wang et al. (2018), Kim & Paik (2019), Hao et al. (2020a), and Oh et al. (2021). Different algorithms are proposed for this problem, and different regret analyses based on various assumptions are provided.

When the unknown parameter becomes a matrix, some recent research have worked on low-rank matrix bandits. For example, Katariya et al. (2017a;b) and Trinh et al. (2020) considered the rank-1 matrix bandit problems, but their results cannot be extended to higher ranks. Jun et al. (2019) studied the bilinear bandit problems where the reward is modelled as a bilinear product of the left arm, the parameter matrix, and the right arm. Kveton et al. (2017) first studied the low-rank matrix bandit, but required strong assumptions on the mean reward matrix. Lu et al. (2021) extended the work by Jun et al. (2019) to a more generalized low-rank matrix bandit problem but their action set was fixed. Hao et al. (2020b) further studied the problem of low-rank tensor bandits.

Despite the recent progress in all the above high dimensional bandit problems, both experimentally and theoretically, prior works are scattered and different algorithms with different assumptions are proposed for these problems. Although in LASSO Bandits, people have already given logarithmic dependence on the high dimension, we are surprised that such dependence has not been shown for other problems and polynomial dependence on the high dimensions still exists. An interesting question here is: does there exist a unified algorithm that works in all the high dimensional bandit problems, and if so, does such an algorithm hold a desirable regret bound in the different settings in terms of the dimensional dependence? In this work, we provide affirmative answers to these questions.

Our work is inspired by the literature on traditional high dimensional statistical analysis (Negahban et al., 2012) and modern high dimensional bandit algorithms (Kim & Paik, 2019; Oh et al., 2021). The only similar prior work we are aware of is the framework by Johnson et al. (2016), but they analyzed a completely different problem setting. Their

¹Department of Statistics, Purdue University, West Lafayette, IN ²Department of Computer Science, Purdue University, West Lafayette, IN. Correspondence to: Wenjie Li <li354@purdue.edu>, Jean Honorio <jhonorio@purdue.edu>.

algorithm is not only much more complicated than ours, but also requires very strong assumptions on the arm sets as well as low-dimensional structural information of the unknown parameter.

In particular, we want to highlight the following contributions of our paper:

- We present a simple and unified algorithm framework named Explore-the-Structure-Then-Commit (ESTC) for high dimensional stochastic bandit problems and provide a problem-independent regret analysis framework for our algorithm. We show that to ensure a desirable regret, one simply needs to ensure that the two special events defined in Section 3 happen with high probability after the exploration stage under mild assumptions.
- We demonstrate the usefulness of our framework by applying it to different high dimensional bandit problems. We show that under a mild unified assumption, our algorithm can achieve desirable regret bounds. We first prove that our algorithm achieves comparable regret bounds in the LASSO bandit problem as a sanity check; we then provide novel bounds in the low-rank matrix bandit and the group sparse matrix bandit which also depend logarithmically on the dimensions. We also show that a simple extension of our algorithm can achieve group sparsity with a desirable regret bound in a new problem: the multi-agent LASSO bandit. We summarize our results and compare them with the most relevant existing works in Table 1.¹
- Our novel framework provides new proofs for different high dimensional bandit problems. Although prior works such as Hao et al. (2020a) has proved similar results in the LASSO Bandit case, the proof techniques for low-rank matrix bandit and group-sparse matrix bandit are highly different and novel, which are also part of the contributions in this paper, e.g., proofs for Lemma 4.5, 4.8, C.2, C.3, and C.4. Interested readers can refer to the Appendix for more details.

2. Preliminaries

In this section, we establish some important preliminary notations and definitions in this paper.

Notations. Given a subspace \mathcal{M} of \mathbb{R}^p , its orthogonal complement \mathcal{M}^\perp is defined as $\mathcal{M}^\perp := \{v \in \mathbb{R}^p \mid \langle u, v \rangle =$

¹For the first regret upper bound by Hao et al. (2020a), i.e. $\mathcal{O}(\tau s^{2/3} T^{2/3} \sqrt{\log(dT)})$, we have contacted the authors personally and confirmed that there is a mistake in the proof of Theorem 4.2 in their paper. The $\sqrt{\log(dT)}$ term is missing in the original theorem. Explanations are provided in Appendix A.

0 for all $u \in \mathcal{M}$ }. The matrix inner product for two matrices A, B of the same size are defined as $\langle\langle A, B \rangle\rangle := \text{trace}(A^T B)$. We use $\|\cdot\|$ for vector norms and $\|\|\cdot\|\|$ for matrix norms. Given a norm $\|\cdot\|$, we use the notation $\theta_{\mathcal{M}}$ to represent the projection of a vector θ onto \mathcal{M} , i.e., $\theta_{\mathcal{M}} = \text{argmin}_{\theta' \in \mathcal{M}} \|\theta - \theta'\|$, and similarly for $\theta_{\mathcal{M}^\perp}$. The dual of a norm $\|\cdot\|$ is defined to be $\|u\|_* := \sup_{\|v\| \leq 1} \langle u, v \rangle$. For the regularization norm R , we denote its dual to be R^* . We use $\tilde{\mathcal{O}}(\cdot)$ to hide the non-dominating factors in big- \mathcal{O} notations. We frequently use the notation $[T]$ for $T \in \mathbb{N}$ to denote the set of integers $\{1, 2, \dots, T\}$. For the reader's reference, a complete list of all the norms and their duals used in this paper is provided in Appendix A.

Multiarmed Contextual Bandits. In modern multiarmed contextual bandit problems, a set of contexts $\{x_{t,a_i}\}_{i=1}^K$ for each arm is generated at every round t , and then the agent chooses an action a_t from the K arms. The contexts are assumed to be sampled i.i.d from a distribution \mathcal{P}_X with respect to t , but the contexts for different arms can be correlated (Chu et al., 2011). After the action is selected, a reward $y_t = f(x_{t,a_t}, \theta^*) + \epsilon_t$ for the chosen action is received, where f is a deterministic function, θ^* is an unknown parameter, and ϵ_t is a zero-mean random noise term which is often assumed to be sub-Gaussian or even conditional sub-Gaussian in a few cases. In this paper, we focus on bandit problems where θ^* is high dimensional but with low-dimensional structure such as a sparse vector, a low-rank matrix, and a group-sparse matrix.

Let $a_t^* = \text{argmax}_{i \in [K]} f(x_{t,a_i}, \theta^*)$ denote the optimal action at each round. We measure the performance of all algorithms by the expectation of the regret, denoted as

$$\mathbb{E}[\text{Regret}(T)] = \mathbb{E} \left[\sum_{t=1}^T f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \right]$$

The goal for all the bandit algorithms studied in this paper is to ensure a sublinear expectation of the regret with respect to T , so that the average regret converges, thus making the chosen actions nearly optimal. Moreover, we require the regret bounds to depend on the low-dimensional structure constraint of the parameter, for example, the number of non-zero elements for a sparse vector, rather than polynomially depend on the high dimension(s), so that the algorithm utilizes the low-dimensional structure of the parameter to reduce the regret.

Optimization Problem. We use the shorthand notations $\mathbf{X}_t = \{x_{i,a_i}\}_{i=1}^t$ and $\mathbf{Y}_t = \{y_i\}_{i=1}^t$ to represent all the contexts of the chosen actions and the rewards received up to time t . Many algorithms designed for multiarmed bandit problems involve solving an online optimization problem with a loss function $L_t(\theta; \mathbf{X}_t, \mathbf{Y}_t)$ and a regularization norm

Table 1. Summary of regret bounds by different research works and their assumptions. **Parameters:** $d, (d_1, d_2)$ are the dimension sizes for vectors and matrices respectively; K denotes the number of arms when the arm set is finite; s denotes the support size of a vector; r denotes the rank of a matrix; τ is a paper-dependent constant. **Arm Set:** “fixed” represents that the arm set is a fixed set \mathcal{A} ; “CBL” represents that different arms have different parameters. “finite” represents that the number of arms is finite; “contextual” represents that K different contexts are generated from a distribution at each round. $\tilde{\mathcal{O}}(\cdot)$ hides the non-dominating terms.

LASSO	REGRET BOUND	ARM SET	ASSUMPTIONS
Lattimore et al. (2015)	$\mathcal{O}(s\sqrt{n})$	fixed	action set is a hypercube
Bastani & Bayati (2020)	$\mathcal{O}(\tau K s^2 \log^2 T)$	CBL	compatibility, margin, optimality
Wang et al. (2018)	$\mathcal{O}(\tau K s^3 \log T)$	CBL	compatibility, margin, optimality
Kim & Paik (2019)	$\mathcal{O}(\tau s \sqrt{T} \log(dT))$	contextual	compatibility
Oh et al. (2021)	$\mathcal{O}(\tau s \sqrt{T} \log(dT))$	contextual	compatibility, symmetry, covariance
Oh et al. (2021)	$\mathcal{O}(\tau \sqrt{sT} \log(dT))$	contextual	RE condition, symmetry, covariance
Hao et al. (2020a)	$\mathcal{O}(s^{2/3} T^{2/3} \sqrt{\log(dT)})$	fixed	action set spans \mathbb{R}^d
Hao et al. (2020a)	$\mathcal{O}(\sqrt{sT} \log(KT))$	finite, fixed	action set spans \mathbb{R}^d , minimum signal
Hao et al. (2020a)	$\Omega(s^{1/3} T^{2/3})$	finite, fixed	action set spans \mathbb{R}^d
This paper (sanity check)	$\mathcal{O}(s^{1/3} T^{2/3} \sqrt{\log(dT)})$	contextual	RE condition
LOW-RANK			
Jun et al. (2019)	$\tilde{\mathcal{O}}((d_1 + d_2)^{3/2} \sqrt{rT})$	fixed	the reward function is bilinear
Johnson et al. (2016)	$\tilde{\mathcal{O}}((d_1 + d_2)^{3/2} \sqrt{rT})$	fixed	prior knowledge of parameter
Lu et al. (2021)	$\tilde{\mathcal{O}}((d_1 + d_2)^{3/2} \sqrt{rT})$	fixed	sub-Gaussian sampling distribution
This paper	$\tilde{\mathcal{O}}(r^{1/3} T^{2/3} \log((d_1 + d_2)))$	contextual	RE condition
GROUP-SPARSE			
Johnson et al. (2016)	$\mathcal{O}(\sqrt{s(d_2 + \log K)} d_1 T)$	fixed	prior knowledge of parameter
This paper	$\tilde{\mathcal{O}}(s^{1/3} \sqrt{d_2} T^{2/3} + s^{1/3} T^{2/3} \sqrt{\log d_1})$	contextual	RE condition
MULTI-AGENT			
This paper	$\mathcal{O}(d_2 s^{1/3} T^{2/3} \sqrt{\log(d_1 T)})$	contextual	RE condition

$R(\theta)$, i.e.,

$$\theta_t \in \operatorname{argmin}_{\theta \in \Theta} \left\{ L_t(\theta; \mathbf{X}_t, \mathbf{Y}_t) + \lambda_t R(\theta) \right\} \quad (1)$$

where λ_t is the regularization parameter chosen differently in different algorithms and Θ is the parameter domain. We also often use the notation $L_t(\theta)$ to represent $L_t(\theta; \mathbf{X}_t, \mathbf{Y}_t)$. The solution θ_t in Eqn. (1) is a regularized estimate of θ^* based on the currently available data. The reason to use the regularization $R(\theta)$ is that it can hopefully identify the low-dimensional structure of θ^* so that θ_t can converge to θ^* fast and the action chosen by the algorithm becomes optimal after a few rounds.

High Dimensional Statistics. Our general theory applies to M -estimators, i.e., the loss has the form $L_t(\theta) = \frac{1}{t} \sum_{\tau=1}^t l_\tau(\theta)$, where each $l_\tau(\theta)$ is a convex and differentiable loss function and $R(\theta)$ is a decomposable norm. M -estimators and decomposable norms are well-accepted concepts in the high dimensional statistics literature (Negahban et al., 2012). We first establish the definition of a decomposable norm.

Definition 2.1. Given a pair of subspaces $\mathcal{M} \subseteq \overline{\mathcal{M}}$, the norm R is said to be decomposable with respect to the subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ iff $R(\theta + \gamma) = R(\theta) + R(\gamma)$, for all $\theta \in \mathcal{M}, \gamma \in \overline{\mathcal{M}}^\perp$.

We present a few examples of the pair of subspaces and the decomposable regularization, each corresponding to one application of our general theory in Section 4.

Example 1. (Sparse Vectors and the l_1 Norm). Let $\theta^* \in \mathbb{R}^d$ be a sparse vector with $s \ll d$ nonzero entries, we denote $S(\theta^*)$ to be the set of non-zero indices of θ^* (i.e., the support). The pair of subspaces $\mathcal{M} \subseteq \overline{\mathcal{M}}$ are chosen as $\mathcal{M} = \overline{\mathcal{M}} = \{\theta \in \mathbb{R}^d \mid \theta_j = 0, \text{ for all } j \notin S(\theta^*)\}$. Then l_1 norm is decomposable with respect to $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ i.e., $\|\theta + \gamma\|_1 = \|\theta\|_1 + \|\gamma\|_1$ for all $\theta \in \mathcal{M}, \gamma \in \overline{\mathcal{M}}^\perp$, since θ and γ have non-zero elements on different entries.

Example 2. (Low Rank Matrices and the Nuclear Norm). Let $\Theta^* \in \mathbb{R}^{d_1 \times d_2}$ be a low rank matrix with rank $r \ll \min\{d_1, d_2\}$. We define the pair of subspace $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ as

$$\begin{aligned} \mathcal{M} &:= \{\Theta \in \mathbb{R}^{d_1 \times d_2} \mid \operatorname{row}(\Theta) \subseteq V, \operatorname{col}(\Theta) \subseteq U\} \\ \overline{\mathcal{M}}^\perp &:= \{\Theta \in \mathbb{R}^{d_1 \times d_2} \mid \operatorname{row}(\Theta) \subseteq V^\perp, \operatorname{col}(\Theta) \subseteq U^\perp\}, \end{aligned}$$

where U and V represent the space of the left and right singular vectors of the target matrix Θ^* . Note that in this case $\mathcal{M} \subsetneq \overline{\mathcal{M}}$. Then the nuclear norm $\|\cdot\|_{nuc}$ is decomposable with respect to $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, i.e., $\|\Theta + \Gamma\|_{nuc} = \|\Theta\|_{nuc} + \|\Gamma\|_{nuc}$ for all $\Theta \in \mathcal{M}, \Gamma \in \overline{\mathcal{M}}^\perp$.

Example 3. (Group-sparse Matrices and the $l_{1,q}$ Norm).

Let $\Theta^* \in \mathbb{R}^{d_1 \times d_2}$ be a matrix with group sparse rows, i.e., each row Θ_i^* is nonzero only if $i \in S(\Theta^*)$, and $|S(\Theta^*)| = s \ll d_1$. Similar to Example 1, we define the pair of subspace $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ as

$$\mathcal{M} = \overline{\mathcal{M}} = \{\Theta \in \mathbb{R}^{d_1 \times d_2} \mid \Theta_i = 0, \text{ for all } i \notin S(\Theta^*)\}$$

The orthogonal complement can be defined with respect to the matrix inner product. Then the $l_{1,q}$ norm, defined as $\|\Theta\|_{1,q} = \sum_{i=1}^{d_1} [\sum_{j=1}^{d_2} |\Theta_{ij}|^q]^{1/q}$, is decomposable with respect to $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$.

Given the decomposable regularization R , it is shown in Negahban et al. (2012) that when the regularization parameter $\lambda_t \geq 2R^*(\nabla L_t(\theta^*))$, the error $\theta_t - \theta^*$ belongs to a constraint set \mathbb{C} defined as follows.

Definition 2.2. For any decomposable R on the subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, the constraint set \mathbb{C} is defined as

$$\mathbb{C} := \left\{ \Delta \mid R(\Delta_{\overline{\mathcal{M}}^\perp}) \leq 3R(\Delta_{\overline{\mathcal{M}}}) + 4R(\theta_{\mathcal{M}^\perp}^*) \right\}$$

Such a set restrains the behavior of the error when the regularization parameter is correctly chosen. Next we present the definition of restricted strong convexity on the loss function $L_t(\theta)$.

Definition 2.3. The loss function $L_t(\theta)$ is said to be restricted strongly convex (RSC) around θ^* with respect to the norm $\|\cdot\|$ with curvature $\alpha > 0$ and tolerance function $Z_t(\theta^*)$ if

$$\begin{aligned} B_t(\theta, \theta^*) &:= L_t(\theta) - L_t(\theta^*) - \langle \nabla L_t(\theta^*), \theta - \theta^* \rangle \\ &\geq \alpha \left\| \theta - \theta^* \right\|^2 - Z_t(\theta^*) \end{aligned}$$

In the high dimensional statistics literature, restricted strong convexity is often ensured by a sufficient number of samples for some specific distributions of \mathbf{X}_t such as the Gaussian distribution (Negahban et al., 2012). In the online case, we need some special assumptions to guarantee that restricted strong convexity holds after a number of rounds. In addition, the following subspace compatibility constant plays a key role in restricting the distance between the true parameter and its estimate by the low dimensional structural constraint, hence generating a desirable regret bound.

Definition 2.4. For a subspace $\overline{\mathcal{M}}$, the subspace compatibility constant with respect to the pair $(\|\cdot\|, R)$ is given by $\phi := \sup_{u \in \overline{\mathcal{M}} \setminus \{0\}} (R(u)/\|u\|)$

For instance, $\phi = \sqrt{s}$ for the $(\|\cdot\|_2, \|\cdot\|_1)$ norm pair and \mathcal{M} defined as in Example 1 because $\sqrt{s}\|u\|_2 \geq \|u\|_1$ for $u \in \mathcal{M}$ by the Cauchy-Schwarz inequality.

3. Main results

In this section, we present our simple and general algorithm, as well as its analysis framework. Let $\alpha > 0$ be a constant to be specified later, we define the following two probability events \mathcal{A}_t and \mathcal{E}_t

$$\begin{aligned} \mathcal{A}_t &:= \left\{ \lambda_t \geq 2R^*(\nabla L_t(\theta^*)) \right\}, \\ \mathcal{E}_t &:= \left\{ B_t(\theta, \theta^*) \geq \alpha \left\| \theta - \theta^* \right\|^2 - Z_t(\theta^*) \right\} \end{aligned}$$

where \mathcal{A}_t represents a correctly chosen regularization parameter and \mathcal{E}_t means that $L_t(\theta)$ is RSC with respect to the undefined norm $\|\cdot\|$ with curvature α and tolerance $Z_t(\theta^*)$ at round t . As we will show below, ensuring that these two events happen with high probability is of vital importance.

3.1. Oracle inequality

We first present a general oracle inequality between the estimate and the true parameter, which is extended from the results in Negahban et al. (2012).

Lemma 3.1. (Oracle Inequality) *If at round t , the probability event \mathcal{A}_t holds and the event \mathcal{E}_t holds for all $\theta - \theta^* \in \mathbb{C}$, then for the solution θ_t of Eqn. (1), the difference $\theta_t - \theta^*$ satisfies the bound*

$$\left\| \theta_t - \theta^* \right\|^2 \leq 9 \frac{\lambda_t^2}{\alpha^2} \phi^2 + \frac{1}{\alpha} \left[2Z_t(\theta^*) + 4\lambda_t R(\theta_{\mathcal{M}^\perp}^*) \right]$$

Remark 3.2. The proof of Lemma 3.1 is provided in Appendix F for completeness. Lemma 3.1 states an oracle inequality for each choice of the pair of norms $(\|\cdot\|, R)$, and the corresponding subspaces where R is decomposable. The difference converges to zero if both λ_t and $Z_t(\theta^*)$ decreases when t increases, and therefore the estimate θ_t becomes more and more accurate as rounds progress.

3.2. The general algorithm

Now we present our Explore-the-Structure-Then-Commit (ESTC) algorithm for high dimensional contextual bandit problems. Our algorithm consists of two stages: the exploration stage where arms are randomly picked and the commit stage where the best arm explored is chosen. Algorithm 1 shows our procedure in detail.

Remark 3.3. Similar algorithms have been proposed in some recent papers. For example, Hao et al. (2020a) proposed an Explore-the-Sparsity-Then-Commit algorithm, but it was only designed and analyzed for fixed action sets and for LASSO bandits. Oh et al. (2021) proposed a sparsity-agnostic algorithm in the LASSO bandit problem, which does not need the sparsity parameter s , but they required much more assumptions than ours in the multi-arm case ($K \geq 3$). Lu et al. (2021) proposed a Low-ESTR algorithm that solved the same optimization problem as our Section

Algorithm 1 Explore-the-Structure-Then-Commit

- 1: **Input:** $\lambda_{T_0}, K \in \mathbb{N}, L_t(\theta), R(\theta), f(x, \theta), \theta_0, T_0$
- 2: Initialize $\mathbf{X}_0, \mathbf{Y}_0 = (\emptyset, \emptyset), \theta_t = \theta_0$
- 3: **for** $t = 1$ to T_0 **do**
- 4: Observe K contexts, $x_{t,1}, x_{t,2}, \dots, x_{t,K}$
- 5: Choose action a_t uniformly randomly
- 6: Receive reward $y_t = f(x_{t,a_t}, \theta^*) + \epsilon_t$
- 7: $\mathbf{X}_t = \mathbf{X}_{t-1} \cup \{x_{t,a_t}\}, \mathbf{Y}_t = \mathbf{Y}_{t-1} \cup \{y_{a_t}\}$
- 8: **end for**
- 9: Compute the estimator θ_{T_0} :

$$\theta_{T_0} \in \operatorname{argmin}_{\theta \in \Theta} \{L_{T_0}(\theta; \mathbf{X}_{T_0}, \mathbf{Y}_{T_0}) + \lambda_{T_0} R(\theta)\}$$

- 10: **for** $t = T_0 + 1$ to T **do**
- 11: Choose action $a_t = \operatorname{argmax}_a f(x_{t,a}, \theta_{T_0})$
- 12: **end for**

4.2 in the low-rank matrix bandit, but they require the lowO-FUL algorithm by Jun et al. (2019) and the algorithm is designed for fixed action sets.

Remark 3.4. Our algorithm generalizes over the prior efforts on different high dimensional bandit problems. The advantages of our algorithm are that **(1)**, it is very simple, **(2)**, it does not require strong assumptions, and **(3)**, it can be applied to different problems while it remains unclear whether prior works can be extended to all the problems we discuss here. Furthermore, we provide comparable regret bounds for existing problems and novel results for new problems (see Table 1, Section 4, and 5).

3.3. Regret analysis

To analyze the regret of Algorithm 1, we impose some weak assumptions on the reward function and the norm of the context. The two assumptions are listed below for the general theorem, and we will demonstrate how they can be easily guaranteed in the specific applications in Section 4.

Assumption 3.5. x is normalized with respect to the norm $\|\cdot\|$, i.e. $\|x\| \leq k_1$ for some constant k_1 .

Assumption 3.6. $f(x, \theta)$ is C_1 -Lipschitz over x and C_2 -Lipschitz over θ with respect to $\|\cdot\|$. i.e.,

$$\begin{aligned} f(x_1, \theta) - f(x_2, \theta) &\leq C_1 \|x_1 - x_2\|, \\ f(x, \theta_1) - f(x, \theta_2) &\leq C_2 \|\theta_1 - \theta_2\| \end{aligned}$$

Remark 3.7. Assumption 3.5 is very standard in the contextual bandit literature, see e.g., Chu et al. (2011); Bastani & Bayati (2020); Kim & Paik (2019); Lu et al. (2021), where the l_2 norm or the Frobenius norm of the contexts is assumed to be either bounded by 1 or a constant x_{max} . In the case where f is normalization-invariant (e.g., the linear case), Assumption 3.5 can be achieved without loss of

generality through normalization of the contexts and the rewards. Assumption 3.6 can also be easily guaranteed with the help of Assumption 3.5 for linear models, as well as for some generalized linear models whose link functions has bounded derivatives such as the logistic model. Given Assumptions 3.5 and 3.6, we present our novel problem-independent theorem of the regret bound.

Theorem 3.8. (Problem Independent Regret Bound) Suppose that Assumption 3.5 and 3.6 hold. Then the expected cumulative regret of Algorithm 1 satisfies the bound

$$\begin{aligned} \mathbb{E}[\operatorname{Regret}(T)] &\leq \underbrace{2C_1 k_1 T_0}_{(a)} + \underbrace{2C_1 k_1 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)]}_{(b)} \\ &+ \underbrace{2C_2 \sum_{t=T_0}^T \sqrt{9 \frac{\lambda_{T_0}^2}{\alpha^2} \phi^2 + \frac{1}{\alpha} [2Z_{T_0}(\theta^*) + 4\lambda_{T_0} R(\theta_{\mathcal{M}^\perp}^*)]}}_{(c)} \end{aligned}$$

Proof Sketch. Instead of assuming a particular form of the model f or the shape of the parameter θ , our proof techniques utilizes the property of all high dimensional bandit problems solved by Algorithm 1 and thus contributes to a general bound. As far as we are aware, no such general analysis have been proved before. The proof starts by showing the instant regret at round t is bounded almost surely.

$$f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \leq 2C_1 k_1$$

which can be used to bound the regret in T_0 and when \mathcal{A}_{T_0} and \mathcal{E}_{T_0} do not happen. Then notice that when $t \geq T_0$, our choice of action in the commit phase implies that $f(x_{t,a_t}, \theta_{T_0}) \geq f(x_{t,a_t^*}, \theta_{T_0})$, we can therefore transform the problem of bounding the regret into a problem of bounding the distance between θ_{T_0} and θ^* by proving that under event \mathcal{A}_{T_0} and \mathcal{E}_{T_0} , for $t \geq T_0$

$$f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \leq 2C_2 \|\theta_{T_0} - \theta^*\|$$

Therefore for any constant $v > 0$, either

$$v \leq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \leq 2C_2 \|\theta_{T_0} - \theta^*\|$$

or $f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \leq v$, which means that either the instantaneous regret is bounded by v , or the distance $\|\theta_{T_0} - \theta^*\|$ is large. The proof is finished by taking v to be the upper bound in Lemma 3.1 multiplied by a constant, and taking expectations and summations over t . \square

Remark 3.9. The full proof is provided in Appendix A. The above regret bound may seem to be complicated at first sight, but we can interpret it in the following way. In the initial T_0 exploration rounds, since we pull arms randomly to collect more samples, we have to consider the worst

case scenario and bound the regret linearly in (a). After enough exploration at T_0 , when the two events \mathcal{A}_{T_0} or \mathcal{E}_{T_0} do not happen, no conclusions can be made on the distance between the estimator θ_{T_0} and the target θ^* , contributing to the second term (b). When both events happen, θ_{T_0} and θ^* are close enough and we can carefully bound the regret with the help of Lemma 3.1, which generates the third term (c). Theorem 3.8 indicates that the expected regret upper bound of Algorithm 1 depends on the probability of \mathcal{A}_{T_0} , \mathcal{E}_{T_0} after a chosen round T_0 , as well as the choice of λ_{T_0} and $Z_{T_0}(\theta^*)$. Therefore, to obtain a sublinear regret, we only need to ensure that the following two things happen in the specific applications.

1. $\mathcal{A}_{T_0}, \mathcal{E}_{T_0}$ are high probability events.
2. $\lambda_{T_0}, Z_{T_0}(\theta^*)$ are carefully chosen so that term (c) is sublinear

For instance, suppose that term (b) is finite, $\lambda_{T_0} = \mathcal{O}(T^{-1/3})$, and $Z_{T_0}(\theta^*) = R(\theta^*_{\mathcal{M}^\perp}) = 0$, then the expected regret is of size $\mathcal{O}(\phi T^{2/3})$ by simple algebra. Such a result is desirable since the regret bound is sublinear with respect to T , and it depends on the subspace compatibility constant instead of the dimension size, so we utilize the low-dimensional structure in θ^* . The final regret bound in Theorem 3.8 will become clearer when we discuss its specific applications.

3.4. High probability events

Finally, we address the two high probability events as they are needed to prove a final regret bound. As we will show in Section 4, the probability of the event \mathcal{A}_t is often decided by the choice of λ_t and the model structure. No further assumptions are needed for \mathcal{A}_t to hold with high probability in all the problems in this paper. However, \mathcal{E}_t does not necessarily hold with high probability even after a large number of rounds, and we need another assumption to guarantee its validity.

Assumption 3.10. (Restricted Eigenvalue Condition)

Let \mathbf{X} denote the matrix where each row is a context vector from an arm. The population Gram matrix $\Sigma = \frac{1}{K} \mathbb{E}[\mathbf{X}^T \mathbf{X}]$ satisfies that there exists some constant $\alpha_0 > 0$ such that $\beta^T \Sigma \beta \geq \alpha_0 \|\beta\|^2$, for all $\beta \in \mathbb{C}$.

Remark 3.11. In the case where the contexts are matrices, Assumption 3.10 represents that the population Gram matrix built from the vectorized contexts satisfies the aforementioned condition. To perform vectorization, we simply need to stack the columns of the matrices and then use the vectorized contexts to obtain the matrix \mathbf{X} . We emphasize that Assumption 3.10 is very general and mild because it is only a requirement on the population, which is satisfied for many distributions, for example, the uniform distribution on

a Euclidean unit ball. Apart from Assumption 3.10, prior works either need many more assumptions such as symmetric distribution and balanced co-variance in (Oh et al., 2021), or very complicated algorithms to make \mathcal{E}_t a high probability event (Kim & Paik, 2019). Our algorithm, however, is simple and general and it can be applied to different problems with only Assumption 3.10.

Remark 3.12. Another popular assumption in the LASSO bandit problem is the compatibility condition (Bastani & Bayati, 2020; Kim & Paik, 2019), which would replace the norm in Assumption 3.10 by the l_1 norm (the regularization norm R in the LASSO bandit problem). Although the compatibility condition is slightly weaker than the restricted eigenvalue condition, we can easily replace the norms in all our arguments by the l_1 norm and still obtain the favorable properties similar to Lemma 3.1. The final regret bound may slightly differ from our results in Section 4.1 in terms of multiplicative constants if we use the compatibility condition, but the proof idea is the same. We use the RE condition in our paper because it is currently unknown whether the compatibility condition can be extended to the matrix bandit case. We leave it as a future work direction.

4. Applications on existing problems

In this section, we present some specific applications of our general framework. Each subsection is organized in the following way. We first clarify all the unspecified notations in the framework, such as the loss, the regularization norm, the compatibility constant, and so on. Then we present a lemma to give the proper choice of the regularization parameter λ_t in each problem. Given the lemma, we derive corollaries of Theorem 3.8 to present the final regret bound of the corresponding algorithm. We emphasize that even though we focus on linear models in all examples for clarity, it is easy to extend our results to nonlinear models that satisfy Assumption 3.5 and 3.6. For example, results for generalized linear models whose link functions are Lipschitz can be easily obtained. Validation experiments that support the correctness of our regret upper bounds are provided in Appendix G

4.1. LASSO bandit

We first consider the LASSO bandit problem. In this case, the reward is assumed to be a linear function of the context of the chosen action $x_{t,a_t} \in \mathbb{R}^d$ and the unknown parameter $\theta^* \in \mathbb{R}^d$, i.e., $y_t = \langle x_{t,a_t}, \theta^* \rangle + \epsilon_t$, where ϵ_t is a (conditional) sub-Gaussian noise. In the LASSO bandit problem, the unknown parameter θ^* is assumed to be sparse with only s non-zero elements and $s \ll d$. This naturally leads to the use of l_1 regularization.

To fit the problem into our framework so that we can get a regret bound, we first clarify the corresponding notations.

The loss function and the regularization are defined to be

$$L_t(\theta) = \frac{1}{2t} \sum_{i=1}^t (y_i - x_{i,a_i}^T \theta)^2, \quad R(\theta) = \|\theta\|_1$$

In this case, we let $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ be defined as in Example 1, then l_1 regularization is decomposable with respect to this pair of spaces. We choose the norm $\|\cdot\|$ in \mathcal{E}_t as the l_2 norm, and thus the compatibility constant $\phi = \sqrt{s}$. Following Chu et al. (2011), we assume that the contexts x_{t,a_i} and the parameter θ^* are all normalized so that $\|x_{t,a_i}\|_2 \leq 1$, $\|\theta^*\|_2 \leq 1$. Therefore Assumption 3.5 and 3.6 are satisfied automatically. Now, we present the following lemma, which provides the choice of λ_t in the LASSO Bandit problem.

Lemma 4.1. *Suppose the noise ϵ_t is conditional σ -sub-Gaussian. For any $\delta \in (0, 1)$, use*

$$\lambda_{T_0} = \frac{2\sigma}{\sqrt{T_0}} \sqrt{2 \log(2d/\delta)}$$

in Algorithm 1, then with probability at least $1 - \delta$, we have $\lambda_{T_0} \geq R^*(\nabla L_{T_0}(\theta^*))$

Given the above lemma, it is easy to apply Theorem 3.8 to get a specific regret bound in the LASSO bandit. Corollary 4.2 follows from taking $\delta = 1/T^2$ in Lemma 4.1 and $T_0 = \Theta(s^{1/3}T^{2/3})$.

Corollary 4.2. *The expected cumulative regret of the Algorithm 1 in the LASSO bandit problem is upper bounded by*

$$\mathbb{E}[\text{Regret}(T)] = \mathcal{O}\left(s^{1/3}T^{2/3} \sqrt{\log(dT)}\right)$$

Remark 4.3. The proofs and the specific algorithm are provided in Appendix B. Note that most of the regret bound depends polynomially on the size of the support s and only logarithmically on the high dimension d . Therefore, Algorithm 1 converges much faster than directly applying linear bandit algorithms such as LinUCB (Chu et al., 2011) to the sparse setting, which satisfies our requirement. Corollary 4.2 indicates that the regret bound is of size $\tilde{\mathcal{O}}(s^{1/3}T^{2/3})$, which matches the regret upper bound proved by Hao et al. (2020a) on the fixed action set and the minimax lower bound up to a logarithmic factor (Hao et al., 2020a).

Remark 4.4. Some recent papers have proved better dependence in terms of the total number of rounds T with more complexly-designed algorithms, e.g. Kim & Paik (2019) or more assumptions, e.g. Oh et al. (2021). Specifically, they have provided $\tilde{\mathcal{O}}(\sqrt{T})$ regret bounds. However, as mentioned by Hao et al. (2020a), these works all need T to be sufficiently large as an additional requirement. In the so-called data-poor regime where T is limited, an $\Omega(T^{2/3})$ regret is unavoidable and thus our regret is optimal (Hao et al., 2020a). Moreover, since vectors are 1-column matrices, such $\Omega(T^{2/3})$ lower bounds are also unavoidable in

the low-rank and group-sparse matrix bandit problems. Our result on LASSO bandit serves as a sanity check of the logarithmic dependence on the dimension and our main goal is to show such dependence in the other problems.

4.2. Low-Rank matrix bandit

Next, we consider the implications of our general theory in the low-rank matrix bandit problems. In this case, the reward is assumed to be a linear function of the low rank matrix $\Theta^* \in \mathbb{R}^{d_1 \times d_2}$ and the context matrix $X_{i,a_i} \in \mathbb{R}^{d_1 \times d_2}$, i.e., $Y_t = \langle X_{t,a_t}, \Theta^* \rangle + \epsilon_t$, where $\text{rank}(\Theta^*) = r \ll \min\{d_1, d_2\}$. We specify the loss function and the regularization norm to be

$$L_t(\Theta) = \frac{1}{2t} \sum_{i=1}^t (y_i - \langle X_{i,a_i}, \Theta \rangle)^2, \quad R(\Theta) = \|\Theta\|_{nuc}$$

By Example 2, the nuclear norm regularization is decomposable on $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ defined by the left and right singular vector spaces of Θ^* . Note that in this case $\phi = \sqrt{2r}$ because the space $\overline{\mathcal{M}}$ contains matrices with rank at most $2r$. The norm we choose in the event \mathcal{E}_t is the Frobenious norm. Also, we assume without loss of generality that the matrices are normalized, so that $\|X_{t,a_i}\|_F \leq 1$, $\|\Theta^*\|_F \leq 1$. Therefore, Assumption 3.5 and 3.6 are automatically satisfied. Similar to Lemma 4.1, the following lemma provides a good choice of λ_t for the low-rank matrix bandit.

Lemma 4.5. *Suppose the noise ϵ_t is σ -sub-Gaussian. At each round t , for any $\delta \in (0, 1)$, use*

$$\lambda_{T_0} = \frac{4\sigma}{\sqrt{T_0}} \sqrt{3 \log(2T_0/\delta) \log^2(2(d_1 + d_2)/\delta)}$$

in Algorithm 1, then with probability at least $1 - \delta$, we have $\lambda_{T_0} \geq R^*(\nabla L_{T_0}(\Theta^*))$

Similarly, we can obtain the regret bound by setting $\delta = 1/T^2$ in Lemma 4.5 and $T_0 = \Theta(r^{1/3}T^{2/3})$.

Corollary 4.6. *The expected cumulative regret of the Algorithm 1 in the low-rank matrix bandit problem is upper bounded by*

$$\mathbb{E}[\text{Regret}(T)] = \mathcal{O}\left(r^{1/3}T^{2/3} \sqrt{\log(T)} \log((d_1 + d_2)T)\right)$$

Remark 4.7. The proofs and the corresponding matrix bandit algorithm are provided in Appendix C. The regret bound in Corollary 4.6 is of size $\tilde{\mathcal{O}}(r^{1/3}T^{2/3} \log((d_1 + d_2)T))$, which polynomially depends on the low rank structure but only logarithmically on the dimensions and is thus better than directly applying linear bandit algorithms to the problem, which is again desirable. We emphasize that our result cannot be directly compared with some prior results, because papers such as Lu et al. (2021); Jun et al. (2019) consider

a fixed set of the context matrices X_{t,a_t} instead of a generating distribution $P_{\mathcal{X}}$, which is different from our setting. Although our dependence on the total number of rounds T is $\tilde{\mathcal{O}}(T^{2/3})$, we have logarithmic dependence on the sum of the dimensions $(d_1 + d_2)$ and thus it is more desirable than bounds that have polynomial dependence on d when the dimensions d_1, d_2 are large.

4.3. Group-sparse matrix bandit

Finally, we also apply our general framework to the group-sparse matrix bandit problem. Similar to the low-rank matrix bandit case, we assume that the expected reward is a linear function of the context matrix $X_{t,a_t} \in \mathbb{R}^{d_1 \times d_2}$ and the optimal parameter $\Theta^* \in \mathbb{R}^{d_1 \times d_2}$, i.e., $Y_t = \langle X_{t,a_t}, \Theta^* \rangle + \epsilon_t$. Define $S(\Theta^*) = \{i \in [d_1] \mid \Theta_i^* = 0\}$ to be the set of non-zero rows in Θ^* and assume that $|S(\Theta^*)| = s \ll d_1$, so the matrix is group sparse. We again specify the sum of squared errors to be the loss, but use the $l_{1,q}$ norm as the regularization norm.

$$L_t(\Theta) = \frac{1}{2t} \sum_{i=1}^t (y_i - \langle X_{i,a_i}, \Theta \rangle)^2, \quad R(\Theta) = \|\Theta\|_{1,q}$$

Note that if $q = 1$, the above problem is equivalent to the LASSO bandit case, because we only need to vectorize all the matrices by stacking their columns. Then $L_t(\Theta)$ becomes the mean squared error and $R(\Theta)$ becomes the l_1 regularization, thus we only consider $q > 1$ here. Define the subspaces using $S(\Theta^*)$ as in Example 3, then the $l_{1,q}$ norm regularization is decomposable.

Similarly, we assume without loss of generality that the matrices are normalized with respect to the Frobenious norm, so that Assumption 3.5 and 3.6 are satisfied automatically. We define the function $\eta(d_2, m) = \max\{1, d_2^m\}$. This notation is used because we need slightly different choices of λ_t for $q \in (1, 2]$ and $q > 2$. Also, the compatibility constant $\phi = \eta(d_2, \frac{1}{q} - \frac{1}{2})\sqrt{s}$, which varies for different q . Similarly, we provide the following lemma for a good choice of the parameter λ_t in group-sparse matrix bandits.

Lemma 4.8. *Suppose that ϵ_t is σ -sub-Gaussian. At each round t , for any $\delta \in (0, 1)$, use*

$$\lambda_{T_0} = \frac{2\sigma}{\sqrt{T_0}} d_2^{1-1/q} + \frac{\sigma \eta(d_2, \frac{1}{2} - \frac{1}{q})}{\sqrt{T_0}} \sqrt{2 \log(2d_1/\delta)}$$

in Algorithm 1, then with probability at least $1 - \delta$, we have $\lambda_{T_0} \geq R^*(\nabla L_{T_0}(\Theta^*))$.

Given the above lemma, we can similarly derive a desirable regret bound.

Corollary 4.9. *With the two functions of d_2*

$$C_1(d_2) = d_2^{1-\frac{1}{q}} \eta(d_2, \frac{1}{q} - \frac{1}{2}),$$

$$C_2(d_2) = \eta(d_2, \frac{1}{q} - \frac{1}{2}) \eta(d_2, \frac{1}{2} - \frac{1}{q})$$

the expected cumulative regret of the Algorithm 1 in the group-sparse matrix bandit problem is upper bounded by

$$\begin{aligned} & \mathbb{E}[\text{Regret}(T)] \\ &= \mathcal{O} \left(C_1(d_2) s^{1/3} T^{2/3} + C_2(d_2) s^{1/3} T^{2/3} \sqrt{\log(d_1 T)} \right) \end{aligned}$$

Remark 4.10. The proofs and the algorithm are provided in Appendix D. The regret bound in Corollary 4.9 depends on the choice of the regularization norm. For example, if $q = 2$, then $R(\Theta)$ is the group-LASSO regularization and the final regret bound is $\mathcal{O}(\sqrt{d_2} s^{1/3} T^{2/3} + s^{1/3} T^{2/3} \sqrt{\log(d_1 T)})$. Besides, the regret bound only has logarithmic dependence on d_1 , so it is desirable since we assume that the matrix is group sparse with $s \ll d_1$. Again, we emphasize that the regret bound is not directly comparable to Johnson et al. (2016) since they consider a fixed action set. Although we have $\tilde{\mathcal{O}}(T^{2/3})$ dependence on T , our regret bound is more desirable when the dimension d_1 is high and our logarithmic dependence on d_1 is preferred.

5. Application on a new problem: the multi-agent LASSO bandit

Apart from direct applications, our general analysis framework can also inspire less-direct use cases. Here we briefly present a novel application where our framework guides a new algorithm and a new regret bound.

Suppose there are d_2 agents solving LASSO bandit problems synchronously, and each agent k receives a stochastic reward $y_t^{(k)} = x_{t,a_t}^{(k)T} \theta^{(k)*} + \epsilon_t^{(k)}$ at every round t . Here, the contexts for each problem are similar and thus the parameters are required to be group-sparse instead of just sparse. The major challenges in this problem are that, first, there are multiple actions and multiple rewards at each round. Second, the Lipschitzness assumption may only be guaranteed for each agent, but not for the whole problem. Therefore, we propose a variant of our algorithm (shown in Algorithm 5 in Appendix E where the agents jointly solve the following optimization problem.

$$\operatorname{argmin} \left\{ \frac{1}{2T_0} \sum_{i=1}^{T_0} \sum_{k=1}^{d_2} (y_i^{(k)} - x_{i,a_i}^{(k)T} \theta^{(k)*})^2 + \lambda_{T_0} \|\Theta\|_{1,2} \right\}$$

Note that the above problem also has the form of Eqn. (1). Therefore, following the same logic in our analysis framework, we prove two high probability events that correspond

to the correctly chosen regularization parameter and the restricted strong convexity in this multi-agent problem, and then an $\mathcal{O}(d_2 s^{1/3} T^{2/3} \sqrt{\log(d_1 T)})$ regret bound is proved in Appendix E.

Theorem 5.1. *Suppose that Assumption 3.5 and Assumption 3.6 hold for each agent (with respect to the $\|\cdot\|_2$ norm). Also suppose that Assumption 3.10 is satisfied. Then the expected cumulative regret of Algorithm 5 is upper bounded by*

$$\mathbb{E}[\text{Regret}(T)] = \mathcal{O}\left(d_2 s^{1/3} T^{2/3} \sqrt{\log 2d_1 T}\right)$$

Remark 5.2. Although the bound in Theorem 5.1 is the same as applying the LASSO bandit algorithm to each agent independently, our regularization can ensure group-sparsity in the parameter, which cannot be guaranteed in the former case. Again, such a regret bound is desirable since we only have logarithmic dependence on d_1 , and $s \ll d_1$.

References

- Abbasi-Yadkori, Y., Pal, D., and Szepesvari, C. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, 2011.
- Alexandra Carpentier, R. M. Bandit theory meets compressed sensing for high dimensional stochastic linear bandit. *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.
- Auer, P. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, pp. 3:397–422, 2002b.
- Bastani, H. and Bayati, M. Online decision-making with high-dimensional covariates. *Operation Research*, pp. 68(1), 276–294, 2020.
- Bobkov, S., Nayar, P., and Tetali, P. Concentration properties of restricted measures with applications to non-lipschitz functions, 2015.
- Buhlmann, P. and Van De Geer, S. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- Chu, W., Li, L., Reyzin, L., and Schapire, R. Contextual bandits with linear payoff functions. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- Hao, B., Lattimore, T., and Wang, M. High-dimensional sparse linear bandits. In *Advances in Neural Information Processing Systems*, 2020a.
- Hao, B., Zhou, J., Wen, Z., and Sun, W. W. Low-rank tensor bandits, 2020b.
- Johnson, N., Sivakumar, V., and Banerjee, A. Structured stochastic linear bandits, 2016.
- Jun, K.-S., Willett, R., Wright, S., and Nowak, R. Bilinear bandits with low-rank structure. In *Proceedings of the 36th International Conference on Machine Learning*, Proceedings of Machine Learning Research, pp. 3163–3172, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- Katariya, S., Kveton, B., Szepesvári, C., Vernade, C., and Wen, Z. Bernoulli rank-1 bandits for click feedback. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 2017a.
- Katariya, S., Kveton, B., Szepesvari, C., Vernade, C., and Wen, Z. Stochastic rank-1 bandits. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017b.
- Kim, G.-S. and Paik, M. C. Doubly-robust lasso bandit. *Advances in Neural Information Processing Systems* 32, pp. 5877–5887, 2019.
- Kveton, B., Szepesvari, C., Rao, A., Wen, Z., Abbasi-Yadkori, Y., and Muthukrishnan, S. Stochastic low-rank bandits, 2017.
- Lattimore, T., Crammer, K., and Szepesvari, C. Linear multi-resource allocation with semi-bandit feedback. In *Advances in Neural Information Processing Systems*, volume 28, pp. 964–972. Curran Associates, Inc., 2015.
- Lu, Y., Meisami, A., and Tewari, A. Low-rank generalized linear bandit problems. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, 2021.
- Negahban, S. N., Ravikumar, P., Wainwright, M. J., and Yu, B. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. *Statistical Science*, pp. 27(4):538–557, 2012.
- Oh, M.-H., Iyengar, G., and Zeevi, A. Sparsity-agnostic lasso bandit. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8271–8280. PMLR, 18–24 Jul 2021.
- Recht, B. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(104):3413–3430, 2011.
- Trinh, C., Kaufmann, E., Vernade, C., and Combes, R. Solving bernoulli rank-one bandits with unimodal thompson sampling. In *Proceedings of the 31st International Conference on Algorithmic Learning Theory*, 2020.

Wang, X., Wei, M., and Yao, T. Minimax concave penalized multi-armed bandit model with high-dimensional covariates. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.

Supplementary Materials to "A Simple Unified Framework for High Dimensional Bandit Problems"

A. Notations and the Main Theorem

A.1. List of Norms

For a vector $v \in \mathbb{R}^d$,

- $\|v\|_2$ is the l_2 norm of v , i.e., $\|v\|_2 = (v_1^2 + v_2^2 + \dots + v_d^2)^{1/2}$, which is self-dual.
- $\|v\|_1$ is the l_1 norm of v , i.e., $\|v\|_1 = |v_1| + |v_2| + \dots + |v_d|$, whose dual norm is the l_∞ norm.
- $\|v\|_\infty$ is the l_∞ norm of v , i.e., $\|v\|_\infty = \max_{i=1,2,\dots,d} |v_i|$, whose dual norm is the l_1 norm.

For a matrix $M \in \mathbb{R}^{d_1 \times d_2}$

- $\|M\|_F$ is the Frobenius norm of M , i.e., $\|M\|_F = (\sum_{i=1}^{d_1} \sum_{j=1}^{d_2} M_{ij}^2)^{1/2}$, which is self-dual.
- $\|M\|_{nuc}$ is the nuclear norm of M , which is defined as $\|M\|_{nuc} = \sum_{i=1}^{\min\{d_1, d_2\}} \sigma_i$, where $\{\sigma_i\}$'s are the singular values of M . The dual norm of the nuclear norm is the l_2 induced operator norm.
- $\|M\|_{op}$ is the operator norm of M induced by the vector norm $\|\cdot\|_2$, i.e., $\|M\|_{op} = \sup_{\|x\|_2=1} \|Mx\|_2$. An important property of $\|\cdot\|_{op}$ is that $\|M\|_{op} = \max_{i=1}^{\min\{d_1, d_2\}} \sigma_i$, where $\{\sigma_i\}$'s are the singular values of M . The dual norm of the induced operator norm is the nuclear norm.
- $\|M\|_{1,q} = \sum_{i=1}^{d_1} [\sum_{j=1}^{d_2} |M_{ij}|^q]^{1/q}$ is the $l_{1,q}$ norm of M . For example if $q = 2$, this corresponds to the group lasso regularization. The dual norm of $l_{1,q}$ norm is the $l_{\infty,p}$ norm, defined as $\|M\|_{\infty,p} = \max_{i=1}^{d_1} [\sum_{j=1}^{d_2} |M_{ij}|^p]^{1/p}$, with the relationship $1/p + 1/q = 1$.
- $\|M\|_{max}$ is the element-wise maximum norm of M , which is defined as $\|M\|_{max} = \max_{i=1}^{d_1} \max_{j=1}^{d_2} |M_{ij}|$.

A.2. Explanations for the Regret Bound by Hao et al. (2020a)

For the results in this subsection, we have personally contacted the authors and confirmed that the logarithmic term is missing in their proof. For clarity, we follow their notations where they have used n instead of T for the total number of rounds and n_1 instead of T_0 for the exploration rounds. The other notations are either the same as ours (e.g. sparsity level of the target vector is denoted by s , dimension size is denoted by d) or just constants (e.g., R_{\max} is the maximum reward, which is 1 in our case). Please refer to the original paper for more detailed notations. Now in their "Section B.3: Proof of Theorem 4.2: regret upper bound", on page 17, the authors have proved that the regret is upper bounded by the following terms in "Step 3: optimize the length of exploration".

With probability at least $1 - \delta$.

$$R_\theta(n) \leq n_1 R_{\max} + (n - n_1) \frac{4}{C_{\min}} \sqrt{\frac{2s^2 (\log(2d) + \log(n_1))}{n_1}} + 3n R_{\max} \exp(-c_1 n_1) \quad (2)$$

In the final step, Hao et al. (2020a) have taken $n_1 = n^{2/3} (s^2 \log(2d))^{1/3} R_{\max}^{-2/3} (2/C_{\min}^2)^{1/3} = \Theta(n^{2/3})$ and compute the final regret bound. However, they have ignored the $\log(n_1)$ term on the numerator of the fractional under the square root sign. In other words, when taking $n_1 = n^{2/3} (s^2 \log(2d))^{1/3} R_{\max}^{-2/3} (2/C_{\min}^2)^{1/3} = \Theta(n^{2/3})$, we get

$$\begin{aligned} R_\theta(n) &\leq n^{2/3} (s^2 \log(2d))^{1/3} (2/C_{\min}^2)^{1/3} R_{\max}^{1/3} + n^{2/3} s^{2/3} R_{\max}^{1/3} \frac{4 \cdot 2^{1/3}}{C_{\min}^{2/3}} \sqrt{\frac{(\log(2d) + \frac{2}{3} \log(n) + \alpha)}{(\log(2d))^{1/3}}} \\ &\quad + 3n R_{\max} \exp(-c_1 n_1) \\ &= \mathcal{O}\left(s^{2/3} n^{2/3} \sqrt{\log(dn)}\right) \end{aligned} \quad (3)$$

where $\alpha = \log\left((s^2 \log(2d))^{1/3} R_{\max}^{-2/3} (2/C_{\min}^2)^{1/3}\right)$ is a constant. Therefore, the regret bound in Hao et al. (2020a) has the same order as our regret bound in Corollary 4.2.

A.3. Proof of Theorem 3.8

Proof. By the Lipschitzness of f over x with respect to $\|\cdot\|$, and the boundedness of $\|x\|$, we have

$$f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \leq C_1 \|x_{t,a_t^*} - x_{t,a_t}\| \leq 2C_1 k_1$$

Then we can decompose the one-step regret from round t into three parts as follows,

$$\begin{aligned} & f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \\ &= [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t \leq T_0) + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) \\ &\quad + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t > T_0, \mathcal{E}_{T_0}^c) \\ &\leq 2C_1 k_1 \mathbb{I}(t \leq T_0) + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) + 2C_1 k_1 \mathbb{I}(t > T_0, \mathcal{E}_{T_0}^c) \\ &= 2C_1 k_1 \mathbb{I}(t \leq T_0) + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t > T_0, f(x_{t,a_t}, \theta_{T_0}) \geq f(x_{t,a_t^*}, \theta_{T_0}), \mathcal{E}_{T_0}) + 2C_1 k_1 \mathbb{I}(t > T_0, \mathcal{E}_{T_0}^c) \end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function. The last equality is due to the choice of $a_t = \operatorname{argmax}_a f(x_{t,a}, \theta_{T_0})$, and thus we know that $f(x_{t,a_t}, \theta_{T_0}) \geq f(x_{t,a_t^*}, \theta_{T_0})$. We focus on the second indicator function now. By the Lipschitzness of f over θ , we have

$$\begin{aligned} & \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) \\ &= \mathbb{I}(t > T_0, f(x_{t,a_t}, \theta_{T_0}) \geq f(x_{t,a_t^*}, \theta_{T_0}), \mathcal{E}_{T_0}) \\ &= \mathbb{I}(t > T_0, f(x_{t,a_t}, \theta_{T_0}) - f(x_{t,a_t^*}, \theta_{T_0}) + f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0}) \\ &= \mathbb{I}(t > T_0, [f(x_{t,a_t}, \theta_{T_0}) - f(x_{t,a_t}, \theta^*)] + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t^*}, \theta_{T_0})] \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0}) \\ &\leq \mathbb{I}(t > T_0, 2C_2 \|\theta_{T_0} - \theta^*\| \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0}) \end{aligned}$$

Substitute the above inequality back and take expectation on both sides of the one-step regret from round t , we get

$$\mathbb{E} [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \leq 2C_1 k_1 \text{ for } t \leq T_0$$

For $t > T_0$ and any constant $v \in \mathbb{R}$, the expectation is bounded by

$$\begin{aligned} & \mathbb{E} [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \\ &\leq \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \mathbb{I}(2C_2 \|\theta_{T_0} - \theta^*\| \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0})] + 2C_1 k_1 \mathbb{P}(\mathcal{E}_{T_0}^c) \\ &\leq \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \mathbb{I}(2C_2 \|\theta_{T_0} - \theta^*\| \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) > 2C_2 v, \mathcal{E}_{T_0})] + 2C_1 k_1 \mathbb{P}(\mathcal{E}_{T_0}^c) \\ &\quad + \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \cdot \\ &\quad \quad \mathbb{I}(2C_2 \|\theta_{T_0} - \theta^*\| \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) < 2C_2 v, \mathcal{E}_{T_0})] \\ &\leq \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \mathbb{I}(2C_2 \|\theta_{T_0} - \theta^*\| \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) > 2C_2 v, \mathcal{E}_{T_0})] \\ &\quad + 2C_1 k_1 \mathbb{P}(\mathcal{E}_{T_0}^c) + 2C_2 v \\ &\leq 2C_1 k_1 \mathbb{P}(\|\theta_{T_0} - \theta^*\| > v, \mathcal{E}_{T_0}) + 2C_1 k_1 \mathbb{P}(\mathcal{E}_{T_0}^c) + 2C_2 v \end{aligned}$$

Now take v to be the upper bound of $\|\theta_{T_0} - \theta^*\|$ in Lemma 3.1, i.e., $v^2 = 9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{1}{\alpha}[2Z_{T_0}(\theta^*) + 4\lambda_{T_0}R(\theta_{\mathcal{M}^\perp}^*)]$. We know by Lemma 3.1 that $\mathbb{P}(\|\theta_{T_0} - \theta^*\| > v, \mathcal{E}_{T_0}) \leq \mathbb{P}(\mathcal{A}_{T_0}^c, \mathcal{E}_{T_0})$. Then the expected cumulative regret becomes

$$\begin{aligned} \mathbb{E}[\text{Regret}(T)] &\leq 2C_1k_1T_0 + \sum_{t=T_0}^T [2C_1k_1\mathbb{P}(\mathcal{A}_{T_0}^c, \mathcal{E}_{T_0}) + 2C_1k_1\mathbb{P}(\mathcal{E}_{T_0}^c) + 2C_2v] \\ &\leq 2C_1k_1T_0 + 2C_1k_1 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2C_2 \sum_{t=T_0}^T \sqrt{9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{1}{\alpha}[2Z_{T_0}(\theta^*) + 4\lambda_{T_0}R(\theta_{\mathcal{M}^\perp}^*)]} \end{aligned}$$

□

A.4. Regret Bound with the Regularization Norm

In this subsection, we provide a regret bound when Assumption 3.5 and 3.6 are assumed base on the regularization norm R , which would correspond to, for example, using l_1 norm and the compatibility condition in LASSO bandit. The regret bound we derive in Corollary A.4 can be easily extended to the LASSO bandit, the low-rank matrix bandit and the group-sparse matrix bandit problems. We first provide the second oracle inequality (Negahban et al., 2012).

Lemma A.1. (Oracle Inequality) *If the probability event \mathcal{A}_t holds and \mathcal{E}_t holds for $\theta - \theta^* \in \mathbb{C} \cap \mathbb{B}(r_t)$, where $r_t^2 \geq 9\frac{\lambda_t^2}{\alpha^2}\phi^2 + \frac{\lambda_t}{\alpha}(2Z_t(\theta^*) + 4R(\theta_{\mathcal{M}^\perp}^*))$. Further suppose that $R(\theta_{\mathcal{M}^\perp}^*) = 0$, then the difference $\theta_t - \theta^*$ satisfies the bound*

$$R(\theta_t - \theta^*) \leq 4\phi \sqrt{9\frac{\lambda_t^2}{\alpha^2}\phi^2 + \frac{2\lambda_t}{\alpha}Z_t(\theta^*)}$$

Assumption A.2. We assume x is normalized with respect to the norm R , i.e., $R(x) \leq k_1$ for some constant k_1 .

Assumption A.3. Here we assume similar Lipschitzness conditions on f and boundedness on $R(x)$, as what we have done in Section 3. We assume $f(x, \theta)$ is C_1 -Lipschitz over x and C_2 -Lipschitz over θ with respect to the norm R . i.e.,

$$\begin{aligned} f(x_1, \theta) - f(x_2, \theta) &\leq C_1R(x_1 - x_2) \\ f(x, \theta_1) - f(x, \theta_2) &\leq C_2R(\theta_1 - \theta_2) \end{aligned}$$

We show that based on such conditions, we can get a similar result as in Theorem 3.8.

Corollary A.4. (Regret Bound) *Let $T_0 \in \mathbb{N}$ be a constant. Suppose that Assumption A.2 and A.3 hold. Also suppose that $\theta^* \in \mathcal{M}$, then the expected cumulative regret satisfies the bound*

$$\mathbb{E}[\text{Regret}(T)] \leq 2C_1k_1T_0 + 2C_1k_1 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 8C_2\phi \sum_{t=T_0}^T \sqrt{9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{2\lambda_{T_0}}{\alpha}Z_{T_0}(\theta^*)}$$

Proof. By the Lipschitzness of f over x with respect to $R(\cdot)$, and the boundedness of $R(x)$, we have

$$f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \leq C_1R(x_{t,a_t^*} - x_{t,a_t}) \leq 2C_1k_1$$

Then we can decompose the one-step regret from round t into three parts as follows

$$\begin{aligned} &f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \\ &= [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t \leq T_0) + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) \\ &\quad + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t \geq T_0, \mathcal{E}_{T_0}^c) \\ &\leq 2C_1k_1\mathbb{I}(t \leq T_0) + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) + 2C_1k_1\mathbb{I}(t > T_0, \mathcal{E}_{T_0}^c) \\ &= 2C_1k_1\mathbb{I}(t \leq T_0) + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \mathbb{I}(t > T_0, f(x_{t,a_t}, \theta_{T_0}) \geq f(x_{t,a_t^*}, \theta_{T_0}), \mathcal{E}_{T_0}) + 2C_1k_1\mathbb{I}(t > T_0, \mathcal{E}_{T_0}^c) \end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function. The last equality is due to the choice of $a_t = \operatorname{argmax}_a f(x_{t,a}, \theta_{T_0})$, and thus we know that $f(x_{t,a_t}, \theta_{T_0}) \geq f(x_{t,a_t^*}, \theta_{T_0})$. We focus on the second indicator function now. By the Lipschitzness of f over θ , we have

$$\begin{aligned} & \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) \\ &= \mathbb{I}(t > T_0, f(x_{t,a_t}, \theta_{T_0}) \geq f(x_{t,a_t^*}, \theta_{T_0}), \mathcal{E}_{T_0}) \\ &= \mathbb{I}(t > T_0, f(x_{t,a_t}, \theta_{T_0}) - f(x_{t,a_t^*}, \theta_{T_0}) + f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0}) \\ &= \mathbb{I}(t > T_0, [f(x_{t,a_t}, \theta_{T_0}) - f(x_{t,a_t}, \theta^*)] + [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t^*}, \theta_{T_0})] \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0}) \\ &\leq \mathbb{I}(t > T_0, 2C_2R(\theta_t - \theta^*) \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0}) \end{aligned}$$

Substitute the above inequality back and take expectation on both sides of the one-step regret from round t , we get

$$\mathbb{E} [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \leq 2C_1k_1 \text{ for } t \leq T_0$$

For $t > T_0$ and any constant $v \in \mathbb{R}$, the expectation is bounded by

$$\begin{aligned} & \mathbb{E} [f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)] \\ &\leq \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \mathbb{I}(2C_2R(\theta_{T_0} - \theta^*) \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), \mathcal{E}_{T_0})] + 2C_1k_1\mathbb{P}(\mathcal{E}_{T_0}^c) \\ &\leq \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \mathbb{I}(2C_2R(\theta_{T_0} - \theta^*) \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) > 2C_2v, \mathcal{E}_{T_0})] + 2C_1k_1\mathbb{P}(\mathcal{E}_{T_0}^c) \\ &+ \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \cdot \\ &\quad \mathbb{I}(2C_2R(\theta_{T_0} - \theta^*) \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*), f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) < 2C_2v, \mathcal{E}_{T_0})] \\ &\leq \mathbb{E} [(f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*)) \mathbb{I}(2C_2R(\theta_{T_0} - \theta^*) \geq f(x_{t,a_t^*}, \theta^*) - f(x_{t,a_t}, \theta^*) > 2C_2v, \mathcal{E}_{T_0})] \\ &\quad + 2C_1k_1\mathbb{P}(\mathcal{E}_{T_0}^c) + 2C_2v \\ &\leq 2C_1k_1\mathbb{P}(R(\theta_{T_0} - \theta^*) > v, \mathcal{E}_{T_0}) + 2C_1k_1\mathbb{P}(\mathcal{E}_{T_0}^c) + 2C_2v \end{aligned}$$

Now take v to be the upper bound of $R(\theta_{T_0} - \theta^*)$ in Lemma A.1, i.e., $v = 4\phi\sqrt{9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{2\lambda_{T_0}}{\alpha}Z_{T_0}(\theta^*)}$. We know by Lemma A.1 that $\mathbb{P}(R(\theta_{T_0} - \theta^*) > v, \mathcal{E}_{T_0}) \leq \mathbb{P}(\mathcal{A}_{T_0}^c, \mathcal{E}_{T_0})$. Then the expected cumulative regret becomes

$$\begin{aligned} \mathbb{E}[\text{Regret}(T)] &\leq 2C_1k_1T_0 + \sum_{t=T_0}^T [2C_1k_1\mathbb{P}(\mathcal{A}_{T_0}^c, \mathcal{E}_{T_0}) + 2C_1k_1\mathbb{P}(\mathcal{E}_{T_0}^c) + 2C_2v] \\ &\leq 2C_1k_1T_0 + 2C_1k_1 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 8C_2\phi \sum_{t=T_0}^T \sqrt{9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{2\lambda_{T_0}}{\alpha}Z_{T_0}(\theta^*)} \end{aligned} \quad \square$$

B. Proof for the LASSO Bandit

B.1. Notations and Algorithm

We first clarify the notations we use throughout the proof for the LASSO bandit. We use the notations $X_t \in \mathbb{R}^{t \times d}$, $Y_t, e_t \in \mathbb{R}^t$ to represent the context matrix, the reward and the error vectors respectively, i.e., $[X_t]_i = x_{i,a_i} \in \mathbb{R}^d$, $[Y_t]_i = y_i$, $[e_t]_i = \epsilon_i, \forall i \in [t]$. The loss function now becomes $L_t(\theta) = \frac{1}{2t}\|Y_t - X_t\theta\|_2^2$. The derivative of $L_t(\theta)$ evaluated at θ^* can be computed as

$$\nabla L_t(\theta^*) = -\frac{1}{t}X_t^T Y_t + X_t^T X_t \theta^* = -\frac{1}{t}X_t^T e_t$$

Algorithm 2 The ESTC Algorithm for LASSO Bandit

- 1: **Input:** $\lambda_{T_0}, K \in \mathbb{N}, L_t(\theta), R(\theta), f(x, \theta), T_0$
- 2: **for** $t = 1$ to T_0 **do**
- 3: Observe K contexts, $x_{t,1}, x_{t,2}, \dots, x_{t,K}$
- 4: Choose action a_t uniformly randomly
- 5: Receive reward $y_t = \langle x_{t,a_t}, \theta^* \rangle + \epsilon_t$
- 6: **end for**
- 7: Compute the estimator θ_{T_0} :

$$\theta_{T_0} \in \operatorname{argmin}_{\theta \in \Theta} \left\{ \frac{1}{2T_0} \sum_{i=1}^{T_0} (y_i - x_{i,a_i}^T \theta)^2 + \lambda_{T_0} \|\theta\|_1 \right\}$$

- 8: **for** $t = T_0 + 1$ to T **do**
 - 9: Choose action $a_t = \operatorname{argmax}_a \langle x_{t,a_t}, \theta_{T_0} \rangle$
 - 10: **end for**
-

The Bregman divergence $B_t(\theta, \theta^*)$ can therefore be computed as

$$\begin{aligned} B_t(\theta, \theta^*) &= L_t(\theta) - L_t(\theta^*) - \langle \nabla L_t(\theta^*), \theta - \theta^* \rangle \\ &= \frac{1}{2t} \|Y_t - X_t \theta\|_2^2 - \frac{1}{2t} \|Y_t - X_t \theta^*\|_2^2 - \frac{1}{t} \langle -X_t^T e_t, \theta - \theta^* \rangle \\ &= \frac{1}{2t} (\theta - \theta^*)^T X_t^T X_t (\theta - \theta^*) \end{aligned}$$

Therefore the event \mathcal{A}_t is equivalent to $\{\lambda_t \geq 2\|\nabla L_t(\theta^*)\|_\infty = \frac{2}{t}\|X_t^T e_t\|_\infty\}$. The event \mathcal{E}_t (RSC condition) is equivalent to

$$\frac{1}{2t} (\theta - \theta^*)^T X_t^T X_t (\theta - \theta^*) \geq \alpha \|\theta - \theta^*\|_2^2 - Z_t(\theta^*)$$

Based on the above Bregman divergence, we define the matrix $\hat{\Sigma}_t \in \mathbb{R}^{d \times d}$ as follows .

$$\hat{\Sigma}_t = \frac{1}{t} \sum_{\tau=1}^t x_{\tau,a_\tau} x_{\tau,a_\tau}^T$$

B.2. Technical Lemmas

The following two Bernstein-type inequalities are very useful in our analysis.

Lemma B.1. (Lemma EC.1 of Bastani & Bayati (2020)) *Let $\{D_k, \mathcal{F}_k\}_{k=1}^\infty$ be a martingale difference sequence, and suppose that D_k is σ -sub-Gaussian in an adapted sense, i.e., for all $\alpha \in \mathbb{R}, \mathbb{E}[e^{\alpha D_k} | \mathcal{F}_{k-1}] \leq e^{\alpha^2 \sigma^2 / 2}$ almost surely. Then, for all $t \geq 0$, we have*

$$\mathbb{P} \left(\left| \sum_{k=1}^n D_k \right| \geq t \right) \leq 2e^{-t^2 / (2n\sigma^2)}$$

Lemma B.2. (Lemma EC.4. of Bastani & Bayati (2020)) *Let x_1, x_2, \dots, x_t be i.i.d. random vectors in \mathbb{R}^d with $\|x_\tau\|_\infty \leq 1$ for all τ . Let $\Sigma = \mathbb{E}[x_\tau x_\tau^T]$ and $\hat{\Sigma}_t = \frac{1}{t} \sum_{\tau=1}^t x_\tau x_\tau^T$. Then for any $w > 0$, we have*

$$\mathbb{P} \left[\left\| \hat{\Sigma}_t - \Sigma \right\|_{\max} \geq 2 \left(w + \sqrt{2w} + \sqrt{\frac{2 \log(d^2 + d)}{t} + \frac{\log(d^2 + d)}{t}} \right) \right] \leq \exp(-tw)$$

B.3. Proof of Lemma 4.1

Proof. Denote $X_t^{(r)}$ to be the r -th column of the matrix X_t . Since we assume that each $\|x_t\|_2 \leq 1$, then the matrix X_t is column normalized, i.e., $\|X_t^{(r)}\|_2 \leq \sqrt{t}$. From the union bound, we know that for any constant $c \in \mathbb{R}$, we have

$$\begin{aligned} \mathbb{P}\left(\frac{2}{t}\|X_t^T e_t\|_\infty \leq c\right) &= \mathbb{P}\left(\forall r \in [d], |\langle e_t, X_t^{(r)} \rangle| \leq \frac{ct}{2}\right) \\ &\geq 1 - \sum_{r=1}^d \mathbb{P}\left(|\langle e_t, X_t^{(r)} \rangle| \geq \frac{ct}{2}\right) \end{aligned}$$

Note that $\{\langle e_t, X_t^{(r)} \rangle\}_{\tau=1}^t$ is a martingale difference sequence adapted to the filtration $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_\tau$, because $\mathbb{E}[\langle e_t, X_t^{(r)} \rangle | \mathcal{F}_{t-1}] = 0$. Also note that each ϵ_t is σ^2 -sub-Gaussian, therefore

$$\begin{aligned} \mathbb{E}\left[\exp\left(\alpha \langle e_t, X_t^{(r)} \rangle\right) | \mathcal{F}_{k-1}\right] &= \mathbb{E}\left[\exp\left(\alpha \sum_{i=1}^t X_{t,i}^{(r)} \epsilon_i\right) | \mathcal{F}_{k-1}\right] \\ &= \mathbb{E}_{X_t^{(r)}}\left[\mathbb{E}_{\epsilon_t | x_\tau}\left[\exp\left(\alpha \sum_{i=1}^t X_{t,i}^{(r)} \epsilon_i\right) | \mathcal{F}_{k-1}, x_\tau\right]\right] \\ &\leq \mathbb{E}_{X_t^{(r)}}\left[\exp\left(\sum_{i=1}^t \alpha^2 X_{t,i}^{2(r)} \sigma^2 / 2\right) | \mathcal{F}_{k-1}\right] \\ &\leq \exp(\alpha^2 \sigma^2 t / 2) \end{aligned}$$

Hence we can use Lemma B.1 and thus

$$\begin{aligned} \mathbb{P}\left(\frac{2}{t}\|X_t^T e_t\|_\infty \leq c\right) &\geq 1 - \sum_{r=1}^d \mathbb{P}\left(|\langle e_t, X_t^{(r)} \rangle| \geq \frac{ct}{2}\right) \\ &\geq 1 - 2de^{-c^2 t / (8\sigma^2)} \end{aligned}$$

If we take $\lambda_t = c = \sqrt{(8\sigma^2 \log \frac{2d}{\delta})/t}$, then with probability $1 - \delta$, we conclude that

$$\mathbb{P}\left(\frac{2}{t}\|X_t^T e_t\|_\infty \leq \lambda_t\right) \geq 1 - \delta \quad \square$$

B.4. Proof of Lemma B.3

Lemma B.3. *Suppose Assumption 3.10 is satisfied, then with probability at least $1 - \exp\left(-\frac{T_0 C(s)}{4}\right)$, we have $B_{T_0}(\theta, \theta^*) \geq \frac{\alpha_0}{4} \|\theta - \theta^*\|_2^2$, for all $T_0 \geq 2 \log(d^2 + d)/C(s)$ and $\theta - \theta^* \in \mathbb{C}$, where $C(s) > 0$ is a constant that depends on s .*

First, we introduce the following two technical lemmas.

Lemma B.4. (**Lemma 9 of Oh et al. (2021)**) *Suppose that $\Sigma_1, \Sigma_2 \in \mathbb{R}^{d \times d}$, $\beta \in \mathbb{C} \cap \mathbb{R}^d$, and that the matrix Σ_1 satisfies the restricted eigenvalue condition $\beta^T \Sigma_1 \beta \geq \alpha \|\beta\|_2^2$ with $\alpha > 0$. Moreover, suppose the two matrices are close enough such that $\|\Sigma_2 - \Sigma_1\|_{\max} \leq \delta$, where $32s\delta \leq \alpha$. Then*

$$\beta^T \Sigma_2 \beta \geq \frac{\alpha}{2} \|\beta\|_2^2$$

Proof. The proof here modifies the proof by Oh et al. (2021), and we provide it for completeness. By the Cauchy-Schwartz

inequality, we have

$$\begin{aligned} |\beta^T \Sigma_1 \beta - \beta^T \Sigma_2 \beta| &= |\beta^T (\Sigma_1 - \Sigma_2) \beta| \\ &\leq \|(\Sigma_1 - \Sigma_2) \beta\|_\infty \|\beta\|_1 \\ &\leq \|\Sigma_1 - \Sigma_2\|_{\max} \|\beta\|_1^2 \\ &\leq \delta \|\beta\|_1^2 \end{aligned}$$

For $\beta \in \mathbb{C}$, we have the inequality $\|\beta_{\overline{M}^c}\|_1 \leq 3\|\beta_{\overline{M}}\|_1$. Thus

$$\|\beta\|_1 \leq 4\|\beta_{\overline{M}}\|_1 \leq 4\phi\|\beta_{\overline{M}}\|_2 \leq 4\phi\|\beta\|_2 \leq 4\phi\sqrt{\frac{1}{\alpha}\beta^T \Sigma_1 \beta}$$

Therefore since we assume that $32s\delta \leq \alpha$, we have

$$|\beta^T \Sigma_1 \beta - \beta^T \Sigma_2 \beta| \leq \delta \|\beta\|_1^2 \leq \frac{16\phi^2\delta}{\alpha} \beta^T \Sigma_1 \beta \leq \frac{1}{2} \beta^T \Sigma_1 \beta$$

By some simple algebra on the above inequality, we know that

$$\beta^T \Sigma_2 \beta \geq \frac{\alpha}{2} \|\beta\|_2^2 \quad \square$$

Lemma B.5. (Distance Between Two Matrices) Define $C(s) = (\sqrt{\frac{\alpha_0}{64s} + 1} - 1)^2 = \Theta(s^{-1})$, where α_0 is defined in Assumption 3.10. Then for all $T_0 \geq 2 \log(d^2 + d)/C(s)$, we have

$$\mathbb{P}\left(\left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{\max} \geq \frac{\alpha_0}{32s}\right) \leq \exp\left(-\frac{T_0 C(s)}{2}\right)$$

Proof. Since we sample uniformly in the exploration stage, we know that the contexts $\{x_{\tau, a_\tau}\}_{\tau=1}^{T_0}$ are i.i.d., and thus for any constant w , by Lemma B.2 we have

$$\mathbb{P}\left(\frac{1}{2}\left\|\hat{\Sigma}_{T_0} - \Sigma\right\|_{\max} \geq w + \sqrt{2w} + \sqrt{\frac{2 \log(d^2 + d)}{T_0}} + \frac{\log(d^2 + d)}{T_0}\right) \leq \exp(-wT_0)$$

Now our choices of w and T_0 are to let both terms concerning w and T_0 to be small, i.e.,

$$w + \sqrt{2w} \leq \frac{\alpha_0}{128s} \quad \text{and} \quad \sqrt{\frac{2 \log(d^2 + d)}{T_0}} + \frac{\log(d^2 + d)}{T_0} \leq \frac{\alpha_0}{128s}$$

Solving the two inequalities leads to

$$w = \frac{1}{2}\left(\sqrt{\frac{\alpha_0}{64s} + 1} - 1\right)^2 \quad \text{and} \quad T_0 \geq 2 \log(d^2 + d) / \left(\sqrt{\frac{\alpha_0}{64s} + 1} - 1\right)^2$$

then we have the following inequality

$$\mathbb{P}\left(\frac{1}{2}\left\|\hat{\Sigma}_{T_0} - \Sigma\right\|_{\max} \geq \frac{\alpha_0}{64s}\right) \leq \exp(-wT_0) = \exp\left(-\frac{T_0 C(s)}{2}\right) \quad \square$$

Finally, we provide the proof of Lemma B.3.

Proof. The proof follows from combining Lemma B.4 and Lemma B.5. Since we have that the restricted eigenvalue condition holds for Σ by Assumption 3.10, i.e.,

$$(\theta - \theta^*)^T \Sigma (\theta - \theta^*) \geq \alpha_0 \|\theta - \theta^*\|_2^2, \forall \theta - \theta^* \in \mathbb{C}$$

Also, the two matrices $\Sigma, \hat{\Sigma}_{T_0}$ are close enough when T_0 is large by Lemma B.5,

$$\mathbb{P} \left(\left\| \Sigma - \hat{\Sigma}_{T_0} \right\|_{\max} \geq \frac{\alpha_0}{32s} \right) \leq \exp \left(-\frac{T_0 C(s)}{2} \right)$$

where $C(s)$ is a constant defined as in Lemma B.5. By Lemma B.4 we can claim that

$$(\theta - \theta^*)^T \hat{\Sigma}_{T_0} (\theta - \theta^*) \geq \frac{\alpha_0}{2} \|\theta - \theta^*\|_2^2, \forall \theta - \theta^* \in \mathbb{C}$$

with high probability $1 - \exp(-\frac{T_0 C(s)}{2})$ when $T_0 \geq 2 \log(d^2 + d)/C(s)$. Therefore

$$B_{T_0}(\theta, \theta^*) = \frac{1}{2} (\theta - \theta^*)^T \hat{\Sigma}_{T_0} (\theta - \theta^*) \geq \frac{\alpha_0}{4} \|\theta - \theta^*\|_2^2, \forall \theta - \theta^* \in \mathbb{C}$$

B.5. Proof of Corollary 4.2

Now given the results in Lemma 4.1, Lemma B.3 and Theorem 3.8, we can easily derive an upper bound for our LASSO bandit algorithm. We specify all the constants here. For example, the Lipschitz constants $C_1 = C_2 = 1$, the bound of the norm $k_1 = 1$, and we set $T_0 \geq 2 \log(d^2 + d)/C(s)$, where $C(s)$ is a constant defined in Lemma B.5. The restricted strong convexity holds with $\alpha = \alpha_0/4$. The compatibility constant $\phi = \sqrt{s}$. Also, $\theta_{\mathcal{M}^\perp}^* = 0$ so $R(\theta_{\mathcal{M}^\perp}^*) = 0$.

Proof. Using Theorem 3.8, we can get the following upper bound

$$\mathbb{E}[\text{Regret}(T)] \leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2 \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} s + \frac{8\lambda_{T_0}}{\alpha_0} Z_{T_0}(\theta^*)}$$

Note that $Z_{T_0}(\theta^*) = 0$. Since we take $\delta = 1/T^2$ in Lemma 4.1, $P(\mathcal{A}_{T_0}^c) \leq 1/T^2$. Furthermore,

$$\sum_{t=T_0}^T \mathbb{P}(\mathcal{A}_{T_0}^c) \leq \sum_{t=T_0}^T \frac{1}{T^2} \leq \sum_{t=1}^T \frac{1}{T^2} \leq \frac{1}{T}$$

Also by our choice of $T_0 = \Theta(s^{1/3} T^{2/3})$, we know that

$$\sum_{t=T_0}^T \mathbb{P}(\mathcal{E}_{T_0}^c) \leq \sum_{t=T_0}^T \exp\left(-\frac{T_0 C(s)}{2}\right) \leq T \exp\left(-\frac{T_0 C(s)}{2}\right) = \mathcal{O}(1)$$

where the last equality is by the fact that $x^{3/2} e^{-ax}$ is bounded above by a constant $(\frac{3}{2a})^3 e^{-\frac{9}{4a}}$ when $a > 0$. Therefore we arrive to the following bound

$$\begin{aligned}
 \mathbb{E}[\text{Regret}(T)] &\leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2 \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} s + \frac{8\lambda_{T_0}}{\alpha_0} Z_{T_0}(\theta^*)} \\
 &\leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 24(T - T_0) \frac{\lambda_{T_0}}{\alpha_0} \sqrt{s} \\
 &\leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + \frac{T}{\sqrt{T_0}} \frac{48\sigma\sqrt{s}}{\alpha_0} \sqrt{2\log(2dT)}
 \end{aligned}$$

Matching the two dominating terms, we get the choice of $T_0 = s^{1/3} (\frac{24\sigma}{\alpha_0})^{2/3} T^{2/3}$ and the regret upper bound

$$\begin{aligned}
 \mathbb{E}[\text{Regret}(T)] &\leq 4s^{1/3} \left(\frac{24\sigma}{\alpha_0}\right)^{2/3} T^{2/3} \sqrt{2\log(2dT)} + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] \\
 &= \mathcal{O}\left(s^{1/3} T^{2/3} \sqrt{\log(dT)}\right)
 \end{aligned}$$

□

C. Proof for the Low-Rank Matrix Bandit

C.1. Notations and Algorithm

Algorithm 3 The ESTC Algorithm for Low Rank Matrix Bandit

- 1: **Input:** $\lambda_{T_0}, K \in \mathbb{N}, L_t(\theta), R(\theta), f(x, \theta), T_0$
- 2: **for** $t = 1$ to T_0 **do**
- 3: Observe K contexts, $X_{t,1}, X_{t,2}, \dots, X_{t,K}$
- 4: Choose action a_t uniformly randomly
- 5: Receive reward $y_t = \langle X_{t,a_t}, \Theta^* \rangle + \epsilon_t$
- 6: **end for**
- 7: Compute the estimator Θ_{T_0} :

$$\Theta_{T_0} \in \operatorname{argmin}_{\Theta \in \Theta} \left\{ \frac{1}{2T_0} \sum_{i=1}^{T_0} (y_i - \langle X_{i,a_i}, \Theta \rangle)^2 + \lambda_t \|\Theta\|_{nuc} \right\}$$

- 8: **for** $t = T_0 + 1$ to T **do**
 - 9: Choose action $a_t = \operatorname{argmax}_a \langle X_{t,a_t}, \Theta_{T_0} \rangle$
 - 10: **end for**
-

We first establish the notations and the corresponding algorithm in low rank matrix bandit problems. We use the shorthand vector notations $\mathcal{X}_i(\Theta^*), Y_t, e_t \in \mathbb{R}^t$ such that $[\mathcal{X}_i(\Theta^*)]_i = \langle X_{i,a_i}, \Theta^* \rangle$, $[Y_t]_i = y_i$, and $[e_t]_i = \epsilon_i$. For any matrix A , $\operatorname{vec}(A)$ denotes the vectorization of the matrix by stacking all rows of the matrix into a vector, i.e., if A has entries

$$\begin{bmatrix} A_{11} & A_{12} & \cdots & A_{1n} \\ A_{21} & A_{22} & \cdots & A_{2n} \\ \cdots & \cdots & \cdots & \cdots \\ A_{m1} & A_{m2} & \cdots & A_{mn} \end{bmatrix}$$

then $\operatorname{vec}(A) = (A_{11}, A_{12}, \dots, A_{1n}, A_{21}, A_{22}, \dots, A_{2n}, \dots, A_{m1}, A_{m2}, \dots, A_{mn})^T$, and $\langle A, B \rangle = \operatorname{tr}(A^T B) = \operatorname{vec}(A)^T \operatorname{vec}(B)$.

The loss function now becomes $L_t(\Theta) = \frac{1}{2t} \|Y_t - \mathcal{X}_t(\Theta)\|_2^2$, and the derivative with respect to Θ^* can be computed as

$$\begin{aligned}\nabla L_t(\Theta^*) &= \frac{\partial}{\partial \Theta^*} \frac{1}{2t} (Y_t - \mathcal{X}_t(\Theta^*))^T (Y_t - \mathcal{X}_t(\Theta^*)) \\ &= \frac{1}{2t} \frac{\partial}{\partial \Theta^*} (-2\mathcal{X}_t(\Theta^*)^T Y_t + \mathcal{X}_t(\Theta^*)^T \mathcal{X}_t(\Theta^*)) \\ &= -\frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i\end{aligned}$$

where the differentiation is based on the chain rule of matrix calculus and the following facts

$$\begin{aligned}\frac{\partial(-2\mathcal{X}_t(\Theta^*)^T Y_t + \mathcal{X}_t(\Theta^*)^T \mathcal{X}_t(\Theta^*))}{\partial \mathcal{X}_t(\Theta^*)} &= -2e_t \\ \frac{\partial[\mathcal{X}_t(\Theta^*)]_i}{\partial \Theta_{j,k}^*} &= \frac{\partial \text{tr}(X_{i,a_i}^T \Theta^*)}{\partial \Theta_{j,k}^*} = X_{i,a_i;j,k}\end{aligned}$$

Therefore the Bregman divergence $B_t(\Theta, \Theta^*)$ in the low rank matrix setting is

$$\begin{aligned}B_t(\Theta, \Theta^*) &= L_t(\Theta) - L_t(\Theta^*) - \langle \nabla L_t(\Theta^*), \Theta - \Theta^* \rangle \\ &= \frac{1}{2t} \|Y_t - \mathcal{X}_t(\Theta)\|_2^2 - \frac{1}{2t} \|Y_t - \mathcal{X}_t(\Theta^*)\|_2^2 + \langle \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i, \Theta - \Theta^* \rangle \\ &= \frac{1}{2t} [-2\mathcal{X}_t(\Theta)^T Y_t + \mathcal{X}_t(\Theta)^T \mathcal{X}_t(\Theta) + 2\mathcal{X}_t(\Theta^*)^T Y_t - \mathcal{X}_t(\Theta^*)^T \mathcal{X}_t(\Theta^*)] + \langle \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i, \Theta - \Theta^* \rangle \\ &= \frac{1}{2t} [-2\mathcal{X}_t(\Theta)^T e_t - 2\mathcal{X}_t(\Theta)^T \mathcal{X}_t(\Theta^*) + \mathcal{X}_t(\Theta)^T \mathcal{X}_t(\Theta) + 2\mathcal{X}_t(\Theta^*)^T e_t + \mathcal{X}_t(\Theta^*)^T \mathcal{X}_t(\Theta^*)] + \langle \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i, \Theta - \Theta^* \rangle \\ &= \frac{1}{2t} [\mathcal{X}_t(\Theta)^T \mathcal{X}_t(\Theta) - 2\mathcal{X}_t(\Theta)^T \mathcal{X}_t(\Theta^*) + \mathcal{X}_t(\Theta^*)^T \mathcal{X}_t(\Theta^*)] \\ &= \frac{1}{2t} [(\mathcal{X}_t(\Theta) - \mathcal{X}_t(\Theta^*))^T (\mathcal{X}_t(\Theta) - \mathcal{X}_t(\Theta^*))] \\ &= \frac{1}{2t} \sum_{i=1}^t \langle X_{i,a_i}, \Theta - \Theta^* \rangle^2\end{aligned}$$

Note that one important property of the above Bregman divergence is

$$\begin{aligned}B_t(\Theta, \Theta^*) &= \frac{1}{2t} \sum_{i=1}^t \langle X_{i,a_i}, \Theta - \Theta^* \rangle^2 \\ &= \frac{1}{2t} \sum_{i=1}^t [\text{vec}(X_{i,a_i})^T \text{vec}(\Theta - \Theta^*)]^2 \\ &= \frac{1}{2t} \sum_{i=1}^t \text{vec}(\Theta - \Theta^*)^T [\text{vec}(X_{i,a_i}) \text{vec}(X_{i,a_i})^T] \text{vec}(\Theta - \Theta^*)\end{aligned}$$

Based on the above discussions, we define the matrix $\hat{\Sigma}_t \in \mathbb{R}^{d_1 d_2 \times d_1 d_2}$ as follows

$$\hat{\Sigma}_t = \frac{1}{t} \sum_{\tau=1}^t \text{vec}(X_{\tau,a_\tau}) \text{vec}(X_{\tau,a_\tau})^T$$

Also by Assumption 3.10, we know that for any $B \in \mathbb{C}$

$$\text{vec}(B)^T \Sigma \text{vec}(B) \geq \alpha_0 \|\text{vec}(B)\|_2^2 = \alpha_0 \|B\|_F^2$$

C.2. Technical Lemmas

The following matrix version of the Bernstein inequality is useful in our analysis of matrix bandits.

Lemma C.1. (Theorem 3.2 of Recht (2011)) *Let X_1, \dots, X_L be independent zero-mean random matrices of dimension $d_1 \times d_2$. Suppose $\rho_k^2 = \max \left\{ \|\mathbb{E}[X_k X_k^T]\|_{op}, \|\mathbb{E}[X_k^T X_k]\|_{op} \right\}$ and $\|X_k\|_{op} \leq M$ almost surely for all k . Then for any $\tau > 0$, we have*

$$\mathbb{P} \left[\left\| \sum_{k=1}^L X_k \right\|_{op} > \tau \right] \leq (d_1 + d_2) \exp \left(\frac{-\tau^2/2}{\sum_{k=1}^L \rho_k^2 + M\tau/3} \right)$$

C.3. Proof of Lemma 4.5

Proof. The requirement on λ_t in terms of the spectral norm of $\nabla L_t(\Theta^*)$ is

$$\mathbb{P} \left(\lambda_t \geq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op} \right) = 1 - \mathbb{P} \left(\lambda_t \leq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op} \right)$$

Here we define the high probability event $\mathcal{B} := \{\max_{i=1, \dots, t} |\epsilon_i| < v\}$, and we first concentrate on the probability of the event conditioned on fixed matrices $\{X_{i,a_i}\}_{i=1}^t$. By the definition of sub-Gaussian random variables.

$$\mathbb{P}(\mathcal{B}^c) = \mathbb{P}(\max_{i=1, \dots, t} |\epsilon_i| > v) \leq \sum_{i=1}^t \mathbb{P}(|\epsilon_i| > v) \leq t \exp\left(-\frac{v^2}{2\sigma^2}\right)$$

where the first inequality is due to the union bound. Hence if we take $v = \sigma \sqrt{2 \log(2t/\delta)}$, then $\mathbb{P}(\mathcal{B}^c) \leq \delta/2$. Under the event \mathcal{B} , the operator norm of each $X_{i,a_i} \epsilon_i$ can be bounded by

$$\|X_{i,a_i} \epsilon_i\|_{op} \leq \|X_{i,a_i}\|_F = \sqrt{\sum_{j=1}^{d_1} \sum_{k=1}^{d_2} X_{a_i;j,k}^2 \epsilon_i^2} \leq v \|X_{i,a_i}\|_F = \sigma \sqrt{2 \log(2t/\delta)}$$

Moreover, under the event \mathcal{B} , we have the following bounds

$$\max \left\{ \|\mathbb{E}[\epsilon_i^2 X_{i,a_i} X_{i,a_i}^T \mid \{X_{i,a_i}\}_{i=1}^t]\|_{op}, \|\mathbb{E}[\epsilon_i^2 X_{i,a_i}^T X_{i,a_i} \mid \{X_{i,a_i}\}_{i=1}^t]\|_{op} \right\} \leq 2\sigma^2 \log(2t/\delta)$$

Therefore we have

$$\begin{aligned} & \mathbb{P} \left(\lambda_t \geq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op} \mid \{X_{i,a_i}\}_{i=1}^t \right) \\ &= 1 - \mathbb{P} \left(\lambda_t \leq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op}, \mathcal{B}^c \mid \{X_{i,a_i}\}_{i=1}^t \right) - \mathbb{P} \left(\lambda_t \leq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op}, \mathcal{B} \mid \{X_{i,a_i}\}_{i=1}^t \right) \\ &\geq 1 - \mathbb{P}(\mathcal{B}^c) - \mathbb{P} \left(\lambda_t \leq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op}, \mathcal{B} \mid \{X_{i,a_i}\}_{i=1}^t \right) \\ &\geq 1 - \left(\frac{\delta}{2} + (d_1 + d_2) \exp\left(\frac{-\lambda_t^2 t^2 / 8}{2t\sigma^2 \log(2t/\delta) + \lambda_t t \sigma \sqrt{2 \log(2t/\delta)} / 6}\right) \right) \end{aligned}$$

where the last inequality is by Lemma C.1 with the corresponding upper bound constants computed above. Now we can further bound the probability using a specific choice of λ_t , i.e., we want that

$$(d_1 + d_2) \exp\left(\frac{-t\lambda_t^2/4}{4\sigma^2 \log(2t/\delta) + \lambda_t \sigma (\sqrt{2 \log(2t/\delta)})/3}\right) \leq \delta/2$$

Reversing the above inequality, we get

$$t\lambda_t^2 \geq 16\sigma^2 \log(2t/\delta) + \frac{4}{3}\lambda_t \sigma \sqrt{2 \log(2t/\delta)} \log\left(\frac{2(d_1 + d_2)}{\delta}\right)$$

Therefore if we take $\lambda_t^2 \geq \frac{128}{9t^2}\sigma^2 \log(2t/\delta) \log^2\left(\frac{2(d_1 + d_2)}{\delta}\right) + \frac{32}{t}\sigma^2 \log(2t/\delta)$, then the above inequality holds. This claim is true because

$$\frac{t\lambda_t^2}{2} \geq 16\sigma^2 \log(2t/\delta) \quad \text{and} \quad \frac{t\lambda_t^2}{2} \geq \frac{4}{3}\lambda_t \sigma \sqrt{2 \log(2t/\delta)} \log\left(\frac{2(d_1 + d_2)}{\delta}\right)$$

Note that $d_1 + d_2 \geq 2$ and hence $\log(2(d_1 + d_2)/\delta) \geq 1$, and thus if we use

$$\lambda_t^2 = \frac{48}{t}\sigma^2 \log(2t/\delta) \log^2(2(d_1 + d_2)/\delta)$$

then we have

$$\begin{aligned} \mathbb{P}\left(\lambda_t \geq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op}\right) &= \mathbb{E}_{\{X_{i,a_i}\}_{i=1}^t} \left[\mathbb{P}\left(\lambda_t \geq 2 \left\| \frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i \right\|_{op} \mid \{X_{i,a_i}\}_{i=1}^t\right) \right] \\ &\geq 1 - \left(\frac{\delta}{2} + \frac{\delta}{2}\right) \geq 1 - \delta \end{aligned} \quad \square$$

C.4. Proof of Lemma C.2

Lemma C.2. *Suppose Assumption 3.10 is satisfied, then with probability at least $1 - \exp\left(-\frac{T_0 C(d_1, d_2)}{4}\right)$, we have $B_{T_0}(\Theta, \Theta^*) \geq \frac{\alpha_0}{4} \|\Theta - \Theta^*\|_F^2$, for all $T_0 \geq 2 \log(d_1 d_2^2 + d_1 d_2)/C(d_1, d_2)$ and $\Theta - \Theta^* \in \mathbb{C}$, where $C(d_1, d_2) > 0$ is a constant that depends on d_1, d_2 .*

First, we introduce the following two technical lemmas.

Lemma C.3. *Suppose that $\Sigma_1, \Sigma_2 \in \mathbb{R}^{d_1 d_2 \times d_1 d_2}$, $\Delta \in \mathbb{C}$, and the matrix Σ_1 satisfies the strongly convex condition $\text{vec}(\Delta)^T \Sigma_1 \text{vec}(\Delta) \geq \alpha \|\Delta\|_F^2$ with $\alpha > 0$. Moreover, suppose the two matrices are close enough such that $\|\Sigma_2 - \Sigma_1\|_{op} \leq \delta$, where $2\delta \leq \alpha$. Then*

$$\text{vec}(\Delta)^T \Sigma_2 \text{vec}(\Delta) \geq \frac{\alpha}{2} \|\Delta\|_F^2$$

Proof. The proof here is inspired by the proof of Lemma 6.17 in [Buhlmann & Van De Geer \(2011\)](#) and Lemma B.4 in Appendix B. A similar argument can be made using the nuclear norms $\|\Delta\|_{nuc}$ since $\|\Delta\|_F \leq \|\Delta\|_{nuc}$. That is, the RSC condition can also be assumed in terms of the nuclear norm. Interested readers may take a deeper look at them by replicating all the lemmas in this section. Note that

$$\begin{aligned} |\text{vec}(\Delta)^T \Sigma_1 \text{vec}(\Delta) - \text{vec}(\Delta)^T \Sigma_2 \text{vec}(\Delta)| &= |\text{vec}(\Delta)^T (\Sigma_1 - \Sigma_2) \text{vec}(\Delta)| \\ &\leq \|(\Sigma_1 - \Sigma_2) \text{vec}(\Delta)\|_2 \|\text{vec}(\Delta)\|_2 \\ &\leq \|\Sigma_1 - \Sigma_2\|_{op} \|\text{vec}(\Delta)\|_2^2 \\ &= \|\Sigma_1 - \Sigma_2\|_{op} \|\Delta\|_F^2 \leq \delta \|\Delta\|_F^2 \end{aligned}$$

where the first inequality is by Cauchy-Schwarz and the second inequality is by the definition of the matrix induced norm $\|\cdot\|$. Therefore since we assume that $2\delta \leq \alpha$, we know that

$$|\text{vec}(\Delta)^T \Sigma_1 \text{vec}(\Delta) - \text{vec}(\Delta)^T \Sigma_2 \text{vec}(\Delta)| \leq \delta \|\Delta\|_F^2 \leq \frac{\delta}{\alpha} \text{vec}(\Delta)^T \Sigma_1 \text{vec}(\Delta) \leq \frac{1}{2} \text{vec}(\Delta)^T \Sigma_1 \text{vec}(\Delta)$$

Therefore

$$\text{vec}(\Delta)^T \Sigma_2 \text{vec}(\Delta) \geq \frac{1}{2} \text{vec}(\Delta)^T \Sigma_1 \text{vec}(\Delta) \geq \frac{\alpha}{2} \|\Delta\|_F^2 \quad \square$$

Lemma C.4. (Distance Between Two Matrices) Define $C(d_1, d_2) = (\sqrt{\frac{\alpha_0}{4\sqrt{d_1 d_2}} + 1} - 1)^2$, where α_0 is defined in Assumption 3.10. Then for all $T_0 \geq 2 \log(d_1^2 d_2^2 + d_1 d_2) / C(d_1, d_2)$, we have

$$\mathbb{P}\left(\left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{op} \geq \frac{\alpha_0}{2}\right) \leq \exp\left(-\frac{T_0 C(d_1, d_2)}{2}\right)$$

Proof. Since we sample uniformly in the exploration stage, we know that the contexts $\{X_{\tau, a_\tau}\}_{\tau=1}^{T_0}$ are i.i.d., and thus for any constant w , by Lemma B.2 we have

$$\mathbb{P}\left(\frac{1}{2} \left\|\hat{\Sigma}_{T_0} - \Sigma\right\|_{\max} \geq w + \sqrt{2w} + \sqrt{\frac{2 \log(d_1^2 d_2^2 + d_1 d_2)}{T_0}} + \frac{\log(d_1^2 d_2^2 + d_1 d_2)}{T_0}\right) \leq \exp(-w T_0)$$

Note that by the property of matrix norms, we have $\left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_2 \leq \sqrt{d_1 d_2} \left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{\max}$, therefore we get

$$\mathbb{P}\left(\frac{1}{2} \left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{\max} \geq \frac{\alpha_0}{4\sqrt{d_1 d_2}}\right) \geq \mathbb{P}\left(\frac{1}{2} \left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{op} \geq \frac{\alpha_0}{4}\right)$$

Now our choices of w and T_0 is to let

$$w + \sqrt{2w} \leq \frac{\alpha_0}{8\sqrt{d_1 d_2}} \quad \text{and} \quad \sqrt{\frac{2 \log(d_1^2 d_2^2 + d_1 d_2)}{T_0}} + \frac{\log(d_1^2 d_2^2 + d_1 d_2)}{T_0} \leq \frac{\alpha_0}{8\sqrt{d_1 d_2}}$$

which leads to the following choices

$$w = \frac{1}{2} \left(\sqrt{\frac{\alpha_0}{4\sqrt{d_1 d_2}} + 1} - 1\right)^2 \quad \text{and} \quad T_0 \geq 2 \log(d_1^2 d_2^2 + d_1 d_2) / \left(\sqrt{\frac{\alpha_0}{4\sqrt{d_1 d_2}} + 1} - 1\right)^2$$

Then we have the following inequality

$$\mathbb{P}\left(\frac{1}{2} \left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{op} \geq \frac{\alpha_0}{4}\right) \leq \mathbb{P}\left(\frac{1}{2} \left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{\max} \geq \frac{\alpha_0}{4\sqrt{d_1 d_2}}\right) \leq \exp(-T_0 w) = \exp\left(-\frac{T_0 C(d_1, d_2)}{2}\right) \quad \square$$

Finally, we provide the proof of Lemma C.2.

Proof. The proof follows from Lemma C.3 and Lemma C.4. Since we have that the restricted eigenvalue condition holds for Σ by Assumption 3.10, i.e.,

$$\text{vec}(\Theta - \Theta^*)^T \Sigma \text{vec}(\Theta - \Theta^*) \geq \alpha_0 \|\Theta - \Theta^*\|_F^2, \forall \Theta - \Theta^* \in \mathbb{C}$$

Also, the two matrices $\Sigma, \hat{\Sigma}_{T_0}$ are close enough when T_0 is large by Lemma C.4,

$$\mathbb{P}\left(\left\|\Sigma - \hat{\Sigma}_{T_0}\right\|_{op} \geq \frac{\alpha_0}{2}\right) \leq \exp\left(-\frac{T_0 C(d_1, d_2)}{2}\right)$$

where $C(d_1, d_2)$ is the constant defined in Lemma C.4. By Lemma C.3 we can claim that

$$\text{vec}(\Theta - \Theta^*)^T \hat{\Sigma}_{T_0} \text{vec}(\Theta - \Theta^*) \geq \frac{\alpha_0}{2} \|\Theta - \Theta^*\|_F^2, \forall \Theta - \Theta^* \in \mathbb{C}$$

with high probability $1 - \exp(-\frac{T_0 C(d_1, d_2)}{2})$ when $T_0 \geq 2 \log(d_1^2 d_2^2 + d_1 d_2) / C(d_1, d_2)$. Therefore

$$B_{T_0}(\Theta, \Theta^*) = \frac{1}{2} (\Theta - \Theta^*)^T \hat{\Sigma}_{T_0} (\Theta - \Theta^*) \geq \frac{\alpha_0}{4} \|\Theta - \Theta^*\|_F^2, \forall \Theta - \Theta^* \in \mathbb{C} \quad \square$$

C.5. Proof of Corollary 4.6

Now given the probability of \mathcal{A}_t and \mathcal{E}_t in Lemma 4.5 and Lemma C.2, we can easily derive a regret upper bound for our low rank matrix bandit algorithm. We specify all the constants here, for example, the Lipschitz constants $C_1 = C_2 = 1$, the bound of the norm $k_1 = 1$, and we set $T_0 \geq 2 \log(d_1^2 d_2^2 + d_1 d_2) / C(d_1, d_2)$, where $C(d_1, d_2)$ is a constant defined in Lemma C.4. The restricted strong convexity holds with $\alpha = \alpha_0 / 4$. The compatibility constant $\phi = \sqrt{2r}$. Also, $\Theta_{\mathcal{M}^\perp}^* = 0$ so $R(\Theta_{\mathcal{M}^\perp}^*) = 0$.

Proof. Using Theorem 3.8, we can get the following upper bound

$$\mathbb{E}[\text{Regret}(T)] \leq 2T_0 + 2C_1 k_1 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2C_2 \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} \phi^2 + \frac{8\lambda_{T_0}}{\alpha_0} Z_{T_0}(\Theta^*)}$$

Since we take $\delta = 1/T^2$ in Lemma 4.5, then $P(\mathcal{A}_t^c) \leq 1/t^2$. Furthermore

$$\sum_{t=T_0}^T \mathbb{P}(\mathcal{A}_{T_0}^c) \leq \sum_{t=T_0}^T \frac{1}{T^2} \leq \sum_{t=1}^T \frac{1}{T^2} \leq \frac{1}{T}$$

Also by our choice of T_0 , we know that

$$\sum_{t=T_0}^T \mathbb{P}(\mathcal{E}_{T_0}^c) \leq \sum_{t=T_0}^T \exp\left(-\frac{T_0 C(d_1, d_2)}{2}\right) \leq T \exp\left(-\frac{T_0 C(d_1, d_2)}{2}\right) = \mathcal{O}(1)$$

where the last equality is by the fact that $x^{3/2} e^{-ax}$ is bounded above by a constant $(\frac{3}{2a})^3 e^{-\frac{9}{4a}}$ when $a > 0$. The last summation term in the regret can be written as

$$\begin{aligned} \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} \phi^2 + \frac{8\lambda_{T_0}}{\alpha_0} Z_{T_0}(\Theta^*)} &\leq \sum_{t=T_0}^T \frac{48\sigma\sqrt{6r}}{\alpha_0} \frac{\sqrt{\log(2T^3)} \log(2(d_1 + d_2)T^2)}{\sqrt{T_0}} \\ &\leq \sum_{t=1}^T \frac{48\sigma\sqrt{6r}}{\alpha_0} \frac{\sqrt{\log(2T^3)} \log(2(d_1 + d_2)T^2)}{\sqrt{T_0}} \\ &\leq (T - T_0) \frac{1}{\sqrt{T_0}} \frac{48\sigma\sqrt{6r}}{\alpha_0} \sqrt{\log(2T^3)} \log(2(d_1 + d_2)T^2) \end{aligned}$$

Therefore we arrive to the following final regret bound

$$\begin{aligned}\mathbb{E}[\text{Regret}(T)] &\leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2 \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} \phi^2 + \frac{8\lambda_{T_0}}{\alpha_0} Z_{T_0}(\Theta^*)} \\ &\leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + \frac{T}{\sqrt{T_0}} \frac{96\sigma\sqrt{6r}}{\alpha_0} \sqrt{\log(2T^3) \log(2(d_1 + d_2)T^2)}\end{aligned}$$

Matching the two dominating terms, we get the choice of $T_0 = r^{1/3}(\frac{48\sqrt{6}\sigma}{\alpha_0})^{2/3}T^{2/3}$ and the regret upper bound

$$\begin{aligned}\mathbb{E}[\text{Regret}(T)] &\leq 4r^{1/3}(\frac{48\sqrt{6}\sigma}{\alpha_0})^{2/3}T^{2/3} \sqrt{2\log(2dT)} \log(2(d_1 + d_2)T^2) + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] \\ &= \mathcal{O}\left(r^{1/3}T^{2/3} \sqrt{\log(T)} \log((d_1 + d_2)T)\right)\end{aligned}$$

□

D. Proof for the Group-sparse Bandit

D.1. Notations and Algorithm

Algorithm 4 The ESTC Algorithm for Group-Sparse Matrix Bandit

- 1: **Input:** $\lambda_{T_0}, K \in \mathbb{N}, L_t(\theta), R(\theta), f(x, \theta), T_0$
- 2: **for** $t = 1$ to T_0 **do**
- 3: Observe K contexts, $X_{t,1}, X_{t,2}, \dots, X_{t,K}$
- 4: Choose action a_t uniformly randomly
- 5: Receive reward $y_t = \langle X_{t,a_t}, \Theta^* \rangle + \epsilon_t$
- 6: **end for**
- 7: Compute the estimator Θ_{T_0} :

$$\Theta_{T_0} \in \operatorname{argmin}_{\Theta \in \Theta} \left\{ \frac{1}{2T_0} \sum_{i=1}^{T_0} (y_i - \langle X_{i,a_i}, \Theta \rangle)^2 + \lambda_t \|\Theta\|_{1,q} \right\}$$

- 8: **for** $t = T_0 + 1$ to T **do**
 - 9: Choose action $a_t = \operatorname{argmax}_a \langle X_{t,a_t}, \Theta_{T_0} \rangle$
 - 10: **end for**
-

We first clarify the notations we use through out this section. As we discuss in Section 4, $\Theta^* = [\theta^{(1)*}, \theta^{(2)*} \dots, \theta^{(d_2)*}]$ is used to denote the matrix whose columns are vectors with similar supports, so that $|S(\Theta^*)| = s \ll d_1$. Similarly, we use the shorthand vector notations $\mathcal{X}_t(\Theta^*), Y_t, e_t \in \mathbb{R}^t$ such that $[\mathcal{X}_t(\Theta^*)]_i = \langle X_{i,a_i}, \Theta^* \rangle$, $[Y_t]_i = y_i$, and $[e_t]_i = \epsilon_i$. By our results in the low-rank matrix bandit problem, the derivative with respect to Θ^* is

$$\begin{aligned}\nabla L_t(\Theta^*) &= \frac{\partial}{\partial \Theta^*} \frac{1}{2t} (Y_t - \mathcal{X}_t(\Theta^*))^T (Y_t - \mathcal{X}_t(\Theta^*)) \\ &= \frac{1}{2t} \frac{\partial}{\partial \Theta^*} (-2\mathcal{X}_t(\Theta^*)^T Y_t + \mathcal{X}_t(\Theta^*)^T \mathcal{X}_t(\Theta^*)) \\ &= -\frac{1}{t} \sum_{i=1}^t X_{i,a_i} \epsilon_i\end{aligned}$$

Therefore the event \mathcal{A}_t is equivalent to

$$\left\{ \lambda_t \geq 2 \|\nabla L_t(\Theta^*)\|_{\infty, q/(q-1)} = 2 \max_{i \in [d_1]} \left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)} \right)^{(q-1)/q} \right\}$$

D.2. Proof of Lemma 4.8

Proof. The proof here basically follows from Lemma 5 in [Negahban et al. \(2012\)](#), except for the function $\eta : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ which we introduce later. From the union bound, we know that for any constant $\lambda_t \in \mathbb{R}$, we have

$$\begin{aligned} \mathbb{P}\left(\lambda_t \geq 2\|\nabla L_t(\Theta^*)\|_{\infty, q/(q-1)}\right) &= \mathbb{P}\left(\forall i \in [d_1], \lambda_t \geq \left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)}\right)^{(q-1)/q}\right) \\ &\geq 1 - \sum_{i=1}^{d_1} \mathbb{P}\left(\lambda_t \leq \left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)}\right)^{(q-1)/q}\right) \end{aligned}$$

Next, we establish a tail bound for the random variable $(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)})^{(q-1)/q}$. Let $X_{a_t}^{(i)}$ be the i -th row of X_{a_t} and define $X = [X_{a_1}^{(i)T}, X_{a_2}^{(i)T}, \dots, X_{a_t}^{(i)T}]$.

For any two σ -sub-Gaussian vectors w, w' , we have

$$\begin{aligned} \left| \frac{1}{t} X w \|_{q/(q-1)} - \frac{1}{t} X w' \|_{q/(q-1)} \right| &\leq \frac{1}{t} \|X(w - w')\|_{q/(q-1)} \\ &= \frac{1}{t} \sup_{\|\theta\|_q=1} \langle X^T \theta, w - w' \rangle \\ &\leq \frac{1}{t} \sup_{\|\theta\|_q=1} \|X^T \theta\|_2 \|w - w'\|_2 \end{aligned}$$

Now if $q \in (1, 2]$, we have

$$\sup_{\|\theta\|_q=1} \|X^T \theta\|_2 \leq \sup_{\|\theta\|_2=1} \|X^T \theta\|_2 \leq \|X^T\|_F \leq \sqrt{t}$$

If $q > 2$, then we have a different inequality

$$\sup_{\|\theta\|_q=1} \|X^T \theta\|_2 \leq d_2^{1/2-1/q} \sup_{\|\theta\|_2=1} \|X^T \theta\|_F \leq d_2^{1/2-1/q} \sqrt{t}$$

Therefore if we define $\eta(d_2, m) = \max\{1, d_2^m\}$, we have $\left| \frac{1}{t} X w \|_{q/(q-1)} - \frac{1}{t} X w' \|_{q/(q-1)} \right| \leq \frac{\eta(d_2, 1/2-1/q)}{\sqrt{t}} \|w - w'\|_2$. Thus the function is Lipschitz with constant $\eta(d_2, 1/2 - 1/q)/\sqrt{t}$. Based on the concentration inequality of measure for Lipschitz functions ([Bobkov et al., 2015](#); [Negahban et al., 2012](#)), we know that

$$\mathbb{P}\left(\frac{1}{t} \|X e_t\|_{q/(q-1)} \geq \mathbb{E}\left[\frac{1}{t} \|X e_t\|_{q/(q-1)}\right] + \sigma \delta\right) \leq 2 \exp\left(-\frac{t \delta^2}{2 \eta(d_2, 1/2 - 1/q)^2}\right).$$

By Lemma 5 in [Negahban et al. \(2012\)](#), we know that the mean is bounded by $2d_2^{1-\frac{1}{q}} \sigma / \sqrt{t}$, therefore

$$\mathbb{P}\left(\frac{1}{t} \|X e_t\|_{q/(q-1)} \geq 2\sigma \frac{d_2^{1-1/q}}{\sqrt{t}} + \sigma \eta(d_2, \frac{1}{2} - \frac{1}{q}) \delta\right) \leq 2 \exp\left(-\frac{t \delta^2}{2}\right).$$

By change of variables and the union bound, we can get the claim in the lemma. \square

D.3. Proof of Corollary 4.9

Now given the probability of \mathcal{A}_t and \mathcal{E}_t in Lemma 4.8 and Lemma C.2, we can derive a regret upper bound for our group-sparse matrix bandit algorithm. Similarly, We specify all the constants here, for example, the Lipschitz constants $C_1 = C_2 = 1$, the bound of the norm $k_1 = 1$, and we set $T_0 \geq 2 \log(d_1^2 d_2^2 + d_1 d_2) / C(d_1, d_2)$, where $C(d_1, d_2)$ is the constant defined in Lemma C.4. The restricted strong convexity holds with $\alpha = \alpha_0 / 4$. The compatibility constant $\phi = \eta(d_2, 1/q - 1/2) \sqrt{s}$. Also, $\Theta_{\mathcal{M}^\perp}^* = 0$ so $R(\Theta_{\mathcal{M}^\perp}^*) = 0$.

Proof. Using Theorem 3.8, we can get the following upper bound

$$\mathbb{E}[\text{Regret}(T)] \leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2 \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} \phi^2 + \frac{8\lambda_{T_0}}{\alpha_0} Z_{T_0}(\Theta^*)}$$

Since we take $\delta = 1/T^2$ in Lemma 4.8, $P(\mathcal{A}_{T_0}^c) \leq 1/T^2$. Furthermore,

$$\sum_{t=T_0}^T \mathbb{P}(\mathcal{A}_{T_0}^c) \leq \sum_{t=T_0}^T \frac{1}{T^2} \leq \sum_{t=1}^T \frac{1}{T^2} \leq 1/T$$

Also by our choice of $T_0 = \Theta(s^{1/3} T^{2/3})$, we know that

$$\sum_{t=T_0}^T \mathbb{P}(\mathcal{E}_{T_0}^c) \leq \sum_{t=T_0}^T \exp\left(-\frac{T_0 C(d_1, d_2)}{2}\right) \leq T \exp\left(-\frac{T_0 C(d_1, d_2)}{2}\right) = \mathcal{O}(1)$$

where the last equality is by the fact that $x^{3/2} e^{-ax}$ is bounded above by a constant $(\frac{3}{2a})^3 e^{-\frac{9}{4a}}$ when $a > 0$. The last summation term in the regret can be written as

$$\begin{aligned} & \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} \phi^2 + \frac{8\lambda_{T_0}}{\alpha_0} Z_t(\Theta^*)} \\ & \leq \sum_{t=T_0}^T \left(\frac{24\sigma\sqrt{s}\eta(d_2, \frac{1}{q} - \frac{1}{2})d_2^{1-\frac{1}{q}}}{\alpha_0} \right) \frac{1}{\sqrt{T_0}} + \sum_{t=T_0}^T \frac{12\sigma\sqrt{s}\eta(d_2, \frac{1}{q} - \frac{1}{2})\eta(d_2, \frac{1}{2} - \frac{1}{q})}{\alpha_0} \frac{\sqrt{2\log(2d_1 T^2)}}{\sqrt{T_0}} \\ & \leq \frac{T - T_0}{\sqrt{T_0}} \left(\frac{24\sigma\sqrt{s}\eta(d_2, \frac{1}{q} - \frac{1}{2})d_2^{1-\frac{1}{q}}}{\alpha_0} + \frac{12\sqrt{s}\sigma\eta(d_2, \frac{1}{q} - \frac{1}{2})\eta(d_2, \frac{1}{2} - \frac{1}{q})}{\alpha_0} \sqrt{2\log(2d_1 T^2)} \right) \end{aligned}$$

Therefore we arrive to the following bound

$$\begin{aligned} \mathbb{E}[\text{Regret}(T)] & \leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2 \sum_{t=T_0}^T \sqrt{\frac{144\lambda_{T_0}^2}{\alpha_0^2} \phi^2 + \frac{8\lambda_{T_0}}{\alpha_0} Z_{T_0}(\Theta^*)} \\ & \leq 2T_0 + 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] \\ & \quad + \frac{T}{\sqrt{T_0}} \left(\frac{24\sqrt{s}\sigma\eta(d_2, \frac{1}{q} - \frac{1}{2})d_2^{1-\frac{1}{q}}}{\alpha_0} + \frac{12\sqrt{s}\sigma\eta(d_2, \frac{1}{q} - \frac{1}{2})\eta(d_2, \frac{1}{2} - \frac{1}{q})}{\alpha_0} \sqrt{2\log(2d_1 T^2)} \right) \end{aligned}$$

Matching the two dominating terms, we get the choice of

$$T_0 = s^{1/3} \left(\frac{12\sigma\eta(d_2, \frac{1}{q} - \frac{1}{2})d_2^{1-\frac{1}{q}}}{\alpha_0} + \frac{6\sigma\eta(d_2, \frac{1}{q} - \frac{1}{2})\eta(d_2, \frac{1}{2} - \frac{1}{q})}{\alpha_0} \sqrt{2\log(2d_1T^2)} \right)^{2/3} T^{2/3} \quad (4)$$

and the following regret upper bound

$$\begin{aligned} \mathbb{E}[\text{Regret}(T)] &\leq 2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] \\ &\quad + 4s^{1/3} \left(\frac{12\sqrt{s}\sigma\eta(d_2, \frac{1}{q} - \frac{1}{2})d_2^{1-\frac{1}{q}}}{\alpha_0} + \frac{6\sqrt{s}\sigma\eta(d_2, \frac{1}{q} - \frac{1}{2})\eta(d_2, \frac{1}{2} - \frac{1}{q})}{\alpha_0} \sqrt{2\log(2d_1T^2)} \right)^{2/3} T^{2/3} \\ &= \mathcal{O} \left(s^{1/3} (C_1(d_2) + C_2(d_2))^{2/3} T^{2/3} (\log(d_1T^2))^{1/3} \right) \end{aligned}$$

where $C_1(d_2) = \eta(d_2, \frac{1}{q} - \frac{1}{2})d_2^{1-\frac{1}{q}}$ and $C_2(d_2) = \eta(d_2, \frac{1}{q} - \frac{1}{2})\eta(d_2, \frac{1}{2} - \frac{1}{q})$. \square

E. Proof for the Multi-agent LASSO

E.1. Notations and Algorithm

Algorithm 5 The ESTC Algorithm for Multiple Agent LASSO Bandit

- 1: **Input:** $\{\lambda_t\}_{t=1}^T, K, d_1, d_2 \in \mathbb{N}, L_t(\theta), R(\theta), f(x, \theta), T_0$
- 2: **for** $t = 1$ to T_0 **do**
- 3: Observe K contexts, $x_{t,1}, x_{t,2}, \dots, x_{t,K}$
- 4: Each Agent k choose action $a_t^{(k)}$ uniformly randomly
- 5: Each Agent receives reward $y_t^{(k)} = f(x_{t,a_t}, \theta^{(k)*}) + \epsilon_t^{(k)}$
- 6: **end for**
- 7: Compute the estimator $\Theta_{T_0} = (\theta_{T_0}^{(1)}, \theta_{T_0}^{(2)}, \dots, \theta_{T_0}^{(d_2)})$:

$$\Theta_{T_0} \in \operatorname{argmin}_{\Theta \in \Theta} \left\{ \frac{1}{2T_0} \sum_{i=1}^{T_0} \sum_{k=1}^{d_2} (y_i^{(k)} - x_{i,a_i}^{(k)T} \theta^{(k)})^2 + \lambda_t \|\Theta\|_{1,q} \right\}$$

- 8: **for** $t = T_0 + 1$ to T **do**
 - 9: Each agent k choose action $a_t^{(k)} = \operatorname{argmax}_a f(x_{t,a}, \theta_{T_0}^{(k)})$
 - 10: **end for**
-

Define $\Theta^* = [\theta^{(1)*}, \theta^{(2)*}, \dots, \theta^{(d_2)*}] \in \mathbb{R}^{d_1 \times d_2}$ and $S(\Theta^*) = \{i \in [d_1] \mid \Theta_i^* = 0\}$ as the set of zero rows, then by our setting $|S(\Theta^*)| = s \ll d_1$. Let the loss be the sum of squared error function and the regularization function be the $l_{1,q}$ norm. More formally,

$$L_t(\Theta) = \frac{1}{2t} \sum_{i=1}^t \sum_{k=1}^{d_2} (y_i^{(k)} - x_{i,a_i}^{(k)T} \theta^{(k)*})^2, R(\Theta) = \|\Theta\|_{1,q}$$

Define $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ as in Example 3, then the $l_{1,q}$ norm is decomposable and $\Theta_{\overline{\mathcal{M}}^\perp}^* = 0$. For each agent j , we use the notations $X_t^{(j)} \in \mathbb{R}^{t \times d_1}, Y_t^{(j)}, e_t^{(j)} \in \mathbb{R}^t$ to represent the context matrix, the reward and the error vectors, i.e., $[X_t^{(j)}]_i = x_{i,a_i}^{(j)}, [Y_t^{(j)}]_i = y_i^{(j)}, [e_t^{(j)}]_i = \epsilon_i^{(j)}, \forall i \in [t], j \in [d_2]$. The derivative with respect to Θ^* can be computed as

$$\nabla_{\Theta^*} L_t(\Theta^*) = \frac{1}{2t} \sum_{i=1}^t \sum_{k=1}^{d_2} \nabla_{\Theta^*} (y_i^{(k)} - x_{i,a_i}^{(k)T} \theta^{(k)*})^2$$

Note that if we take partial derivatives, we get

$$\begin{aligned}\frac{\partial L_t(\Theta^*)}{\partial \theta^{(j)*}} &= \frac{1}{2t} \sum_{i=1}^t \sum_{k=1}^{d_2} \frac{\partial}{\partial \theta^{(j)*}} (y_i^{(k)} - x_{i,a_i}^{(k)T} \theta^{(k)*})^2 \\ &= \frac{1}{2t} \sum_{i=1}^t \frac{\partial}{\partial \theta^{(j)*}} (y_i^{(j)} - x_{i,a_i}^{(j)T} \theta^{(j)*})^2 \\ &= -\frac{1}{t} X_t^{(j)T} e_t^{(j)}\end{aligned}$$

Therefore $\nabla L_t(\Theta^*) = -\frac{1}{t} [X_t^{(1)T} e_t^{(1)}, X_t^{(2)T} e_t^{(2)}, \dots, X_t^{(d_2)T} e_t^{(d_2)}]$. Now we can compute the Bregman divergence as follows.

$$\begin{aligned}B_t(\Theta, \Theta^*) &= L_t(\Theta) - L_t(\Theta^*) - \langle \nabla L_t(\Theta^*), \Theta - \Theta^* \rangle \\ &= \frac{1}{2t} \sum_{k=1}^{d_2} \|Y_t^{(k)} - X_t^{(k)} \theta^{(k)}\|_2^2 - \frac{1}{2t} \sum_{k=1}^{d_2} \|Y_t^{(k)} - X_t^{(k)} \theta^{(k)*}\|_2^2 + \frac{1}{t} \sum_{k=1}^{d_2} e_t^{(k)T} X_t^{(k)} (\theta^{(k)} - \theta^{(k)*}) \\ &= \frac{1}{2t} \sum_{k=1}^{d_2} (\theta^{(k)} - \theta^{(k)*})^T X_t^{(k)T} X_t^{(k)} (\theta^{(k)} - \theta^{(k)*})\end{aligned}$$

Therefore the event \mathcal{A}_t is equivalent to

$$\left\{ \lambda_t \geq 2 \|\nabla L_t(\Theta^*)\|_{\infty, q/(q-1)} = 2 \max_{i \in [d_1]} \left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)} \right)^{(q-1)/q} \right\}$$

The event \mathcal{E}_t (RSC condition) is equivalent to

$$\left\{ \frac{1}{2t} \sum_{k=1}^{d_2} (\theta^{(k)} - \theta^{(k)*})^T X_t^{(k)T} X_t^{(k)} (\theta^{(k)} - \theta^{(k)*}) \geq \alpha \|\Theta - \Theta^*\|_F^2 - Z_t(\Theta^*) \right\}$$

E.2. Useful Lemmas

Lemma E.1. (Good Choice of Lambda) Suppose that ϵ_t is σ -sub-Gaussian. For any $\delta \in (0, 1)$, use

$$\lambda_t = 2\sigma \frac{d_2^{1-1/q}}{\sqrt{t}} + \sigma\eta \left(d_2, \frac{1}{2} - \frac{1}{q} \right) \sqrt{\frac{2 \log(2d_1/\delta)}{t}}$$

at each round t in Algorithm 5, then with probability at least $1 - \delta$, we have $\lambda_t \geq R^*(\nabla L_t(\Theta^*))$.

Proof. From the union bound, we know that for any constant $\lambda_t \in \mathbb{R}$, we have

$$\begin{aligned}\mathbb{P} \left(\lambda_t \geq 2 \|\nabla L_t(\Theta^*)\|_{\infty, q/(q-1)} \right) &= \mathbb{P} \left(\forall i \in [d_1], \lambda_t \geq \left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)} \right)^{(q-1)/q} \right) \\ &\geq 1 - \sum_{i=1}^{d_1} \mathbb{P} \left(\lambda_t \leq \left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)} \right)^{(q-1)/q} \right)\end{aligned}$$

Note that the j -th row of $\nabla L_t(\Theta^*)$ consists of $[\sum_i (x_{i,a_i}^{(1)})_j \epsilon_i^{(1)}, \sum_i (x_{i,a_i}^{(2)})_j \epsilon_i^{(2)}, \dots, \sum_i (x_{i,a_i}^{(d_2)})_j \epsilon_i^{(d_2)}]$, which are all $\sqrt{t}\sigma$ -sub-Gaussian random variables.

For any two $\sqrt{t}\sigma$ -sub-Gaussian vectors w, w' , we have

$$\begin{aligned} \left| \left\| \frac{1}{t}w \right\|_{q/(q-1)} - \left\| \frac{1}{t}w' \right\|_{q/(q-1)} \right| &\leq \frac{1}{t} \|(w - w')\|_{q/(q-1)} \\ &= \frac{1}{t} \sup_{\|\theta\|_q=1} \langle \theta, w - w' \rangle \\ &\leq \frac{1}{t} \sup_{\|\theta\|_q=1} \|\theta\|_2 \|w - w'\|_2 \end{aligned}$$

Now if $q \in (1, 2]$, we have $\|\theta\|_2 \leq \|\theta\|_q = 1$. If $q > 2$, then we have a different inequality $\|\theta\|_2 \leq d_2^{1/2-1/q} \|\theta\|_q \leq d_2^{1/2-1/q}$. Therefore if we define $\eta(d_2, m) = \max\{1, d_2^m\}$, we have $\left| \left\| \frac{1}{t}w \right\|_{q/(q-1)} - \left\| \frac{1}{t}w' \right\|_{q/(q-1)} \right| \leq \frac{\eta(d_2, 1/2-1/q)}{t} \|w - w'\|_2$. Thus the function is Lipschitz with constant $\eta(d_2, 1/2 - 1/q)/t$. Based on the concentration inequality of measure for Lipschitz functions (Bobkov et al., 2015; Negahban et al., 2012), we know that

$$\mathbb{P} \left(\left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)} \right)^{(q-1)/q} \geq \mathbb{E} \left[\left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)} \right)^{(q-1)/q} \right] + \delta \right) \leq 2 \exp \left(- \frac{t^2 \delta^2}{2\eta(d_2, 1/2 - 1/q)^2 t \sigma^2} \right).$$

By similar arguments as in the proof of Lemma 4.8, we know that

$$\mathbb{P} \left(\left(\sum_{j=1}^{d_2} |\nabla L_t(\Theta^*)_{ij}|^{q/(q-1)} \right)^{(q-1)/q} \geq 2\sigma \frac{d_2^{1-1/q}}{\sqrt{t}} + \sigma \eta \left(d_2, \frac{1}{2} - \frac{1}{q} \right) \delta \right) \leq 2 \exp \left(- \frac{t \delta^2}{2} \right).$$

By change of variables and the union bound, we can get the results in the lemma. \square

Lemma E.2. (RSC condition) Suppose Assumption 3.10 is satisfied, then with probability at least $1 - d_2 \exp \left(- \frac{T_0 C(s)}{2} \right)$, we have $B_{T_0}(\Theta, \Theta^*) \geq \frac{\alpha_0}{4} \|\Theta - \Theta^*\|_F^2$, for all $T_0 \geq 2 \log(2d_1^2 + d_1)/C(s)$ and $\Theta - \Theta^* \in \mathbb{C}$, where $C(s)$ is a constant defined in Lemma B.5.

Proof. Note that the RSC condition (event \mathcal{E}_t) is equivalent to

$$\begin{aligned} \frac{1}{2t} \sum_{k=1}^{d_2} (\theta^{(k)} - \theta^{(k)*})^T X_t^{(k)T} X_t^{(k)} (\theta^{(k)} - \theta^{(k)*}) &\geq \alpha \|\Theta - \Theta^*\|_F^2 - Z_t(\Theta^*) \\ &= \sum_{k=1}^{d_2} \alpha \|\theta^{(k)} - \theta^{(k)*}\|_2^2 - Z_t(\Theta^*) \end{aligned}$$

By the results in Lemma B.3 in the LASSO bandit problem, we know that with high probability $1 - \exp \left(- \frac{T_0 C(s)}{2} \right)$

$$\frac{1}{2} (\theta^{(k)} - \theta^{(k)*})^T \hat{\Sigma}_{T_0}^{(k)} (\theta^{(k)} - \theta^{(k)*}) \geq \frac{\alpha_0}{4} \|\theta^{(k)} - \theta^{(k)*}\|_2^2$$

when $T_0 \geq 2 \log(d_1^2 + d_1)/C(s)$, where $C(s)$ is defined in Lemma B.5. Therefore by the Frechet inequality, we know that

$$\begin{aligned}
 & P \left(\sum_{k=1}^{d_2} \frac{1}{2t} (\theta^{(k)} - \theta^{(k)*})^T X_{T_0}^{(k)T} X_{T_0}^{(k)} (\theta^{(k)} - \theta^{(k)*}) \geq \sum_{k=1}^{d_2} \frac{\alpha_0}{2} \|\theta^{(k)} - \theta^{(k)*}\|_2^2 \right) \\
 & \geq P \left(\frac{1}{2t} (\theta^{(k)} - \theta^{(k)*})^T X_{T_0}^{(k)T} X_{T_0}^{(k)} (\theta^{(k)} - \theta^{(k)*}) \geq \frac{\alpha_0}{2} \|\theta^{(k)} - \theta^{(k)*}\|_2^2, \forall k \in [d_2] \right) \\
 & \geq \sum_{k=1}^{d_2} P \left(\frac{1}{2t} (\theta^{(k)} - \theta^{(k)*})^T X_{T_0}^{(k)T} X_{T_0}^{(k)} (\theta^{(k)} - \theta^{(k)*}) \geq \frac{\alpha_0}{2} \|\theta^{(k)} - \theta^{(k)*}\|_2^2 \right) - (d_2 - 1) \\
 & \geq 1 - d_2 \exp \left(-\frac{T_0 C(s)}{2} \right)
 \end{aligned} \quad \square$$

E.3. Proof of Theorem 5.1

Proof. By the boundedness of $\|x\|_2$ and $\|\theta\|_2$, we know that

$$f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) = x_{t,a_t^*}^{(k)T} \theta^{(k)*} - x_{t,a_t}^{(k)T} \theta^{(k)*} \leq (x_{t,a_t^*}^{(k)T} - x_{t,a_t}^{(k)T}) \theta^{(k)*} \leq 2$$

Then we can decompose the one-step regret from round t across different agents into three parts as follows

$$\begin{aligned}
 & \sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) \\
 & = \left[\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) \right] \mathbb{I}(t \leq T_0) + \left[\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) \right] \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) \\
 & \quad + \left[\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) \right] \mathbb{I}(t \geq T_0, \mathcal{E}_{T_0}^c) \\
 & \leq 2d_2 \mathbb{I}(t \leq T_0) + \left[\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) \right] \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) + 2d_2 \mathbb{I}(t > T_0, \mathcal{E}_{T_0}^c) \\
 & = 2d_2 \mathbb{I}(t \leq T_0) + \left[\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) \right] \mathbb{I}(t > T_0, f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)}) \geq f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)}), \forall k \in [d_2], \mathcal{E}_{T_0}) \\
 & \quad + 2d_2 \mathbb{I}(t > T_0, \mathcal{E}_{T_0}^c)
 \end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function. The last equality is due to the choice of $a_t^{(k)} = \operatorname{argmax}_a f(x_{t,a}^{(k)}, \theta_{T_0}^{(k)})$, and thus we know that $f(x_{t,a}^{(k)}, \theta_{T_0}^{(k)}) \geq f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)})$. We focus on the second indicator function now. By the Lipschitzness of f over θ , we

have

$$\begin{aligned}
 & \mathbb{I}(t > T_0, \mathcal{E}_{T_0}) \\
 &= \mathbb{I}\left(t > T_0, f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)}) \geq f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)}), \forall k \in [d_2], \mathcal{E}_{T_0}\right) \\
 &= \mathbb{I}\left(t > T_0, f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)}) - f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)}) + f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*})\right. \\
 &\quad \left. \geq f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*}), \forall k, \mathcal{E}_{T_0}\right) \\
 &= \mathbb{I}\left(t > T_0, [f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*})] + [f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*})]\right. \\
 &\quad \left. \geq f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*}), \forall k \in [d_2], \mathcal{E}_{T_0}\right) \\
 &\leq \mathbb{I}\left(t > T_0, 2\|\theta_{T_0}^{(k)} - \theta^{(k)*}\|_2 \geq f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*}), \forall k \in [d_2], \mathcal{E}_{T_0}\right) \\
 &\leq \mathbb{I}\left(t > T_0, 2\sum_{k=1}^{d_2} \|\theta_{T_0}^{(k)} - \theta^{(k)*}\|_2 \geq \sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*}), \mathcal{E}_{T_0}\right) \\
 &\leq \mathbb{I}\left(t > T_0, 2\sqrt{d_2}\|\Theta_{T_0} - \Theta^*\|_F \geq \sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta_{T_0}^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta_{T_0}^{(k)*}), \mathcal{E}_{T_0}\right)
 \end{aligned}$$

where the last inequality is by the Cauchy-Schwarz Inequality. Substitute the above inequality back and take expectation on both sides of the one-step regret from round t , we get

$$\mathbb{E}\left[\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*})\right] \leq 2d_2 \text{ for } t \leq T_0$$

For $t > T_0$ and any constant $v \in \mathbb{R}$, the expectation is bounded by

$$\begin{aligned}
 & \mathbb{E}\left[\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*})\right] \\
 &\leq \mathbb{E}\left[\left(\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*})\right) \mathbb{I}\left(2\sqrt{d_2}\|\Theta_{T_0} - \Theta^*\|_F \geq \sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}), \mathcal{E}_{T_0}\right)\right] \\
 &\quad + 2d_2\mathbb{P}(\mathcal{E}_{T_0}^c) \\
 &\leq \mathbb{E}\left[\left(\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*})\right) \mathbb{I}\left(2\sqrt{d_2}\|\Theta_{T_0} - \Theta^*\|_F \geq \sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*}) \geq 2\sqrt{d_2}v, \mathcal{E}_{T_0}\right)\right] + 2d_2\mathbb{P}(\mathcal{E}_{T_0}^c) + 2\sqrt{d_2}v \\
 &\leq \mathbb{E}\left[\left(\sum_{k=1}^{d_2} f(x_{t,a_t^*}^{(k)}, \theta^{(k)*}) - f(x_{t,a_t}^{(k)}, \theta^{(k)*})\right) \mathbb{I}(\|\Theta_{T_0} - \Theta^*\|_F \geq v), \mathcal{E}_{T_0}\right] + 2d_2\mathbb{P}(\mathcal{E}_{T_0}^c) + 2\sqrt{d_2}v \\
 &\leq 2d_2\mathbb{P}(\|\Theta_{T_0} - \Theta^*\|_F \geq v, \mathcal{E}_{T_0}) + 2d_2\mathbb{P}(\mathcal{E}_{T_0}^c) + 2\sqrt{d_2}v
 \end{aligned}$$

By Lemma E.2, the RSC condition is satisfied when $T_0 \geq 2\log(2d_1^2 + d_1)/C(s)$, where $C(s)$ is defined in Lemma B.5.

Now take v to be the upper bound of $\|\Theta_{T_0} - \Theta^*\|_F$ in Lemma 3.1, i.e., $v^2 = 9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{1}{\alpha}[2Z_{T_0}(\theta^*) + 4\lambda_{T_0}R(\theta_{\mathcal{M}^\perp}^*)]$. We

know by Lemma 3.1, the expected cumulative regret becomes

$$\begin{aligned} \mathbb{E}[\text{Regret}(T)] &\leq 2d_2T_0 + \sum_{t=T_0}^T \left[2d_2\mathbb{P}(\mathcal{A}_{T_0}^c, \mathcal{E}_{T_0}) + 2d_2\mathbb{P}(\mathcal{E}_{T_0}^c) + 2\sqrt{d_2}v \right] \\ &\leq 2d_2T_0 + 2d_2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2\sqrt{d_2} \sum_{t=1}^T \sqrt{9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{\lambda_{T_0}}{\alpha} [2Z_{T_0}(\theta^*) + 4R(\theta_{\mathcal{M}^\perp}^*)]} \end{aligned}$$

By setting $\delta = 1/T^2$ in Lemma E.1 and Lemma E.2, the second term can be bounded as

$$2d_2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_t^c) + \mathbb{P}(\mathcal{E}_t^c)] \leq \sum_{t=1}^T \left[\frac{1}{T^2} + d_2 \exp\left(-\frac{T_0 C(s)}{2}\right) \right] = \mathcal{O}(1)$$

The last term can be bounded as

$$\begin{aligned} &2\sqrt{d_2} \sum_{t=1}^T \sqrt{9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{1}{\alpha} [2Z_{T_0}(\theta^*) + 4\lambda_{T_0}R(\theta_{\mathcal{M}^\perp}^*)]} \\ &\leq 2\sqrt{d_2} \sum_{t=1}^T \sqrt{144\frac{\lambda_{T_0}^2}{\alpha_0^2}\phi^2 + \frac{8\lambda_t}{\alpha_0} Z_{T_0}(\theta^*)} \\ &\leq 2\sqrt{d_2\frac{144\phi^2}{\alpha_0^2}} \sum_{t=1}^T 2 \left(\sigma \frac{d_2^{1-1/q}}{\sqrt{T_0}} + \sigma\eta(d_2, \frac{1}{2} - \frac{1}{q}) \sqrt{\frac{2\log 2d_1T^2}{T_0}} \right) \\ &\leq 2\eta(d_2, \frac{1}{q} - \frac{1}{2}) \sqrt{\frac{144}{\alpha_0^2} d_2 s} \sum_{t=1}^T 2 \left(\sigma \frac{d_2^{1-1/q}}{\sqrt{T_0}} + \sigma\eta(d_2, \frac{1}{2} - \frac{1}{q}) \sqrt{\frac{2\log 2d_1T^2}{T_0}} \right) \\ &\leq \frac{48}{\alpha_0} \eta(d_2, \frac{1}{q} - \frac{1}{2}) \sqrt{d_2 s} \frac{T}{\sqrt{T_0}} \left(\sigma d_2^{1-1/q} + \sigma\eta(d_2, \frac{1}{2} - \frac{1}{q}) \sqrt{2\log 2d_1T^2} \right) \end{aligned}$$

Take $q = 2$ and match the two dominating terms, we get the choice of $T_0 = \Theta(d_2 s^{1/3} T^{2/3})$, Therefore the final regret bound is of size

$$\begin{aligned} &2d_2T_0 + 2d_2 \sum_{t=T_0}^T [\mathbb{P}(\mathcal{A}_{T_0}^c) + \mathbb{P}(\mathcal{E}_{T_0}^c)] + 2\sqrt{d_2} \sum_{t=1}^T \sqrt{9\frac{\lambda_{T_0}^2}{\alpha^2}\phi^2 + \frac{1}{\alpha} [2Z_{T_0}(\theta^*) + 4\lambda_{T_0}R(\theta_{\mathcal{M}^\perp}^*)]} \\ &= \mathcal{O}\left(d_2 s^{1/3} T^{2/3} \sqrt{\log 2d_1 T}\right) \end{aligned} \quad \square$$

F. Proofs Related to the Oracle Inequalities

Recall the definition of the constraint set \mathbb{C} . We define the subset with bounded norm $\mathbb{K}(\delta) := \mathbb{C} \cap \{\Delta | \Delta \leq \delta\}$. Also define the function $F_t(\Delta)$ as

$$F_t(\Delta) = L_t(\theta^* + \Delta) - L_t(\theta^*) + \lambda_t(R(\theta^* + \Delta) - R(\theta^*))$$

F.1. Proof of Lemma 3.1

First, we introduce the following two lemmas from [Negahban et al. \(2012\)](#)

Lemma F.1. (Lemma 3 of Negahban et al. (2012)) For any vectors θ^* , Δ , and a decomposable norm R on \mathcal{M} , $\overline{\mathcal{M}}^\perp$, we have the following inequality

$$R(\theta^* + \Delta) - R(\theta^*) \geq R(\Delta_{\overline{\mathcal{M}}^\perp}) - R(\Delta_{\overline{\mathcal{M}}}) - 2R(\theta^*_{\mathcal{M}^\perp})$$

Lemma F.2. (Lemma 4 of Negahban et al. (2012)) If $F_t(\Delta) > 0$ for all vectors $\Delta \in \mathbb{K}(\delta)$ for a constant δ , then $\|\theta_t - \theta^*\| \leq \delta$

Now, we provide the proof of Lemma 3.1

Proof. Note that $\delta^2 \leq r_t^2$, by the restricted strong convexity of $L_t(\theta^*)$ on $\mathbb{K}(\delta)$, we have

$$\begin{aligned} F_t(\Delta) &= L_t(\theta^* + \Delta) - L_t(\theta^*) + \lambda_t(R(\theta^* + \Delta) - R(\theta^*)) \\ &\geq \langle \nabla L_t(\theta^*), \Delta \rangle + \alpha \|\Delta\|^2 - Z_t(\theta^*) + \lambda_t(R(\theta^* + \Delta) - R(\theta^*)) \\ &\geq \langle \nabla L_t(\theta^*), \Delta \rangle + \alpha \|\Delta\|^2 - Z_t(\theta^*) + \lambda_t(R(\Delta_{\overline{\mathcal{M}}^\perp}) - R(\Delta_{\overline{\mathcal{M}}}) - 2R(\theta^*_{\mathcal{M}^\perp})) \\ &\geq -|\langle \nabla L_t(\theta^*), \Delta \rangle| + \alpha \|\Delta\|^2 - Z_t(\theta^*) + \lambda_t(R(\Delta_{\overline{\mathcal{M}}^\perp}) - R(\Delta_{\overline{\mathcal{M}}}) - 2R(\theta^*_{\mathcal{M}^\perp})) \\ &\geq -R^*(\nabla L_t(\theta^*))R(\Delta) + \alpha \|\Delta\|^2 - Z_t(\theta^*) + \lambda_t(R(\Delta_{\overline{\mathcal{M}}^\perp}) - R(\Delta_{\overline{\mathcal{M}}}) - 2R(\theta^*_{\mathcal{M}^\perp})) \\ &\geq \alpha \|\Delta\|^2 - Z_t(\theta^*) + \lambda_t(R(\Delta_{\overline{\mathcal{M}}^\perp}) - R(\Delta_{\overline{\mathcal{M}}}) - 2R(\theta^*_{\mathcal{M}^\perp})) - \frac{\lambda_t}{2}R(\Delta) \\ &\geq \alpha \|\Delta\|^2 - Z_t(\theta^*) + \lambda_t(R(\Delta_{\overline{\mathcal{M}}^\perp}) - R(\Delta_{\overline{\mathcal{M}}}) - 2R(\theta^*_{\mathcal{M}^\perp})) - \frac{\lambda_t}{2}R(\Delta_{\overline{\mathcal{M}}^\perp}) - \frac{\lambda_t}{2}R(\Delta_{\overline{\mathcal{M}}}) \\ &= \alpha \|\Delta\|^2 - Z_t(\theta^*) + \frac{\lambda_t}{2}(R(\Delta_{\overline{\mathcal{M}}^\perp}) - 3R(\Delta_{\overline{\mathcal{M}}}) - 4R(\theta^*_{\mathcal{M}^\perp})) \end{aligned}$$

where the second inequality is by Lemma F.1. The fourth inequality is by the generalized Cauchy Schwarz inequality. The fifth inequality is because of our setting $\lambda_t \geq 2R(\nabla L_t(\theta^*))$. The last inequality is because of the triangle inequality $R(\Delta) = R(\Delta_{\overline{\mathcal{M}}^\perp} + \Delta_{\overline{\mathcal{M}}}) \leq R(\Delta_{\overline{\mathcal{M}}^\perp}) + R(\Delta_{\overline{\mathcal{M}}})$. Now by the subspace compatibility constant, we know that

$$R(\Delta_{\overline{\mathcal{M}}}) \leq \phi \|\Delta_{\overline{\mathcal{M}}}\| = \phi \|\Pi_{\overline{\mathcal{M}}}(\Delta) - \Pi_{\overline{\mathcal{M}}}(0)\| \leq \phi \|\Delta - 0\|$$

where the last inequality is because $0 \in \overline{\mathcal{M}}$ and because the projection operation is non-expansive. Therefore we can continue to lower bound $F_t(\Delta)$ in the following way.

$$\begin{aligned} F_t(\Delta) &\geq \alpha \|\Delta\|^2 - Z_t(\theta^*) - \frac{\lambda_t}{2}(3R(\Delta_{\overline{\mathcal{M}}}) + 4R(\theta^*_{\mathcal{M}^\perp})) \\ &\geq \alpha \|\Delta\|^2 - Z_t(\theta^*) - \frac{3\lambda_t\phi}{2}\|\Delta\| - 2\lambda_t R(\theta^*_{\mathcal{M}^\perp}) \end{aligned}$$

Now since we take $\|\Delta\|^2 = \delta^2 = 9\frac{\lambda_t^2}{\alpha^2}\phi^2 + \frac{1}{\alpha}[2Z_t(\theta^*) + 4\lambda_t R(\theta^*_{\mathcal{M}^\perp})]$, by the same algebraic manipulations in Negahban et al. (2012), we have $F_t(\Delta) > 0$. Now by Lemma F.2, we know that

$$\|\theta_t - \theta^*\|^2 \leq 9\frac{\lambda_t^2}{\alpha^2}\phi^2 + \frac{1}{\alpha}[2Z_t(\theta^*) + 4\lambda_t R(\theta^*_{\mathcal{M}^\perp})]$$

The second oracle inequality in Lemma A.1 can be proved easily by the triangle inequality decomposition and the definition of \mathbb{C} . That is

$$\begin{aligned} R(\theta_t - \theta^*) &= R((\theta_t - \theta^*)_{\overline{\mathcal{M}}^\perp} + (\theta_t - \theta^*)_{\overline{\mathcal{M}}}) \leq R((\theta_t - \theta^*)_{\overline{\mathcal{M}}^\perp}) + R((\theta_t - \theta^*)_{\overline{\mathcal{M}}}) \\ &\leq 4R((\theta_t - \theta^*)_{\overline{\mathcal{M}}}) + 4R(\theta^*_{\mathcal{M}^\perp}) \\ &\leq 4\phi \|(\theta_t - \theta^*)_{\overline{\mathcal{M}}}\| + 4R(\theta^*_{\mathcal{M}^\perp}) \\ &\leq 4\phi \|\theta_t - \theta^*\| + 4R(\theta^*_{\mathcal{M}^\perp}) \end{aligned}$$

If $\theta^* \in \mathcal{M}$, then $R(\theta_{\mathcal{M}^\perp}^*) = 0$. Therefore we know that

$$R(\theta_t - \theta^*) \leq 4\phi \sqrt{9 \frac{\lambda_t^2}{\alpha^2} \phi^2 + \frac{2}{\alpha} Z_t(\theta^*)}$$

□

G. Validation Experiments

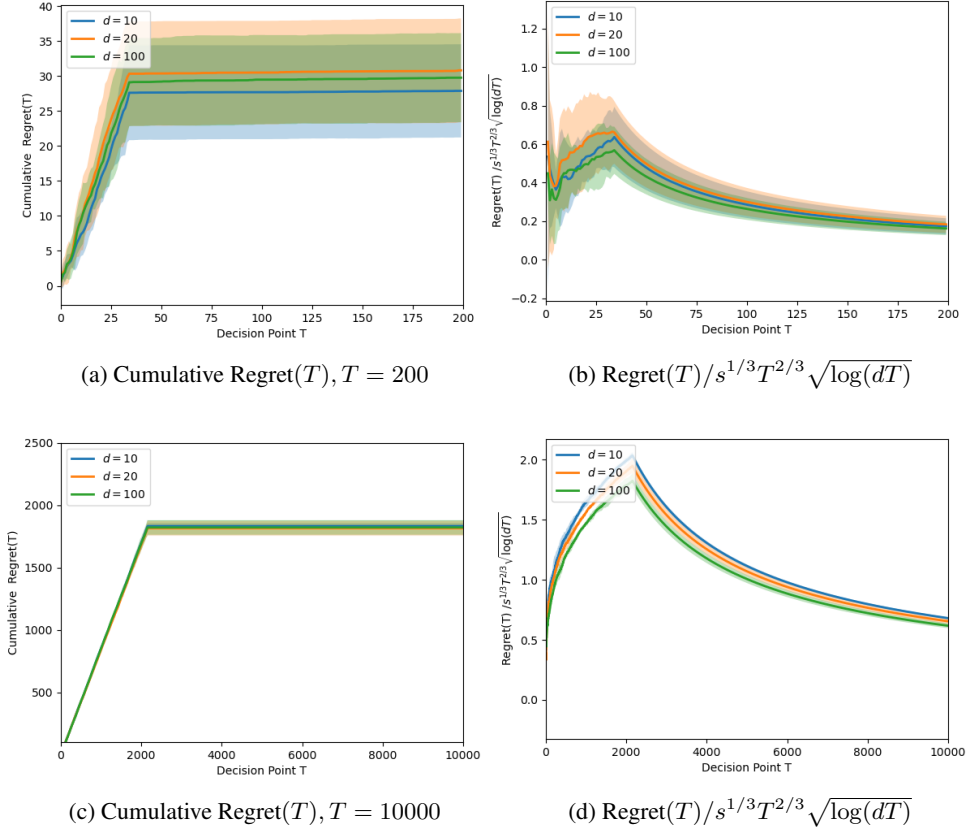


Figure 1. The experimental results in LASSO Bandit problem when $T = 200$ and $T = 10000$

In this section, we provide some experiments in order to validate our claims in the theoretical results. We choose to run our algorithms in the different high dimensional bandit problems and validate the corresponding regret upper bounds in Corollary 4.2, Corollary 4.6, Corollary 4.9 and Theorem 5.1. For each problem, we plot the cumulative regret $R(T)$ as well as $R(T)/B(T)$, where $B(T)$ is the upper bound derived in the specific application. For each setting in each problem, we run the algorithm for 10 independent runs and plot the mean results with one standard deviation error bars.

- **LASSO Bandit.** We generate our true θ^* by randomly choosing its non-zero indices, and then generate each of its non-zero values uniformly randomly from $[0,1]$ and then perform the normalization. We set $K = 10$ so that there are ten different contexts available at each round. The contexts $\{x_{t,a_i}\}_{i=1}^K$ are generated from the zero-mean and identity-covariance normal distribution. We choose three different dimension sizes $d = 10, 20, 100$. The sparsity level is chosen to be $s = 5$ in all three settings. The experiment results are shown in Figure 1. As we can observe, the figures show that the regret is at most a constant times $s^{1/3}T^{2/3}\sqrt{\log(dT)}$, which only depends on the dimension d logarithmically and thus satisfies our requirement. Therefore, our bound correctly delineate the order of the regret.
- **Low-rank Matrix Bandit.** We first randomly generate r vectors of dimension d_2 , and then we generate our true Θ^* by randomly choosing its rows from the generated r vectors, and thus the matrix is low-rank. We set $K = 10$ so that

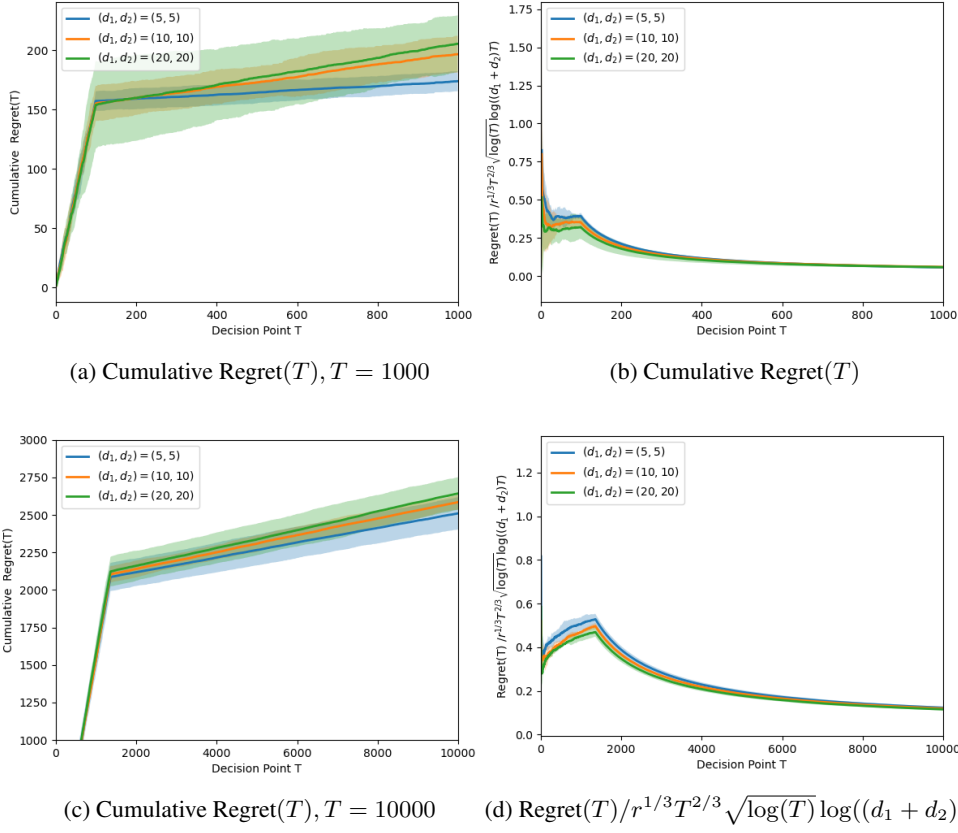


Figure 2. The experimental results in Low-rank Matrix Bandit problem when $T = 1000$ and $T = 10000$

there are ten different contexts available at each round. The contexts $\{X_{t,a_i}\}_{i=1}^K$ are generated from the zero-mean and identity-covariance normal distribution. We choose three different dimension sizes $(d_1, d_2) = (5, 5), (10, 10), (20, 20)$. The rank is chosen to be $r = 2$ in all three settings. The experiment results are shown in Figure 2. As we can observe, the figures show that the regret is at most a constant times $r^{1/3}T^{2/3}\sqrt{\log(T)}\log((d_1 + d_2)T)$, which only depends on the dimensions $d_1 + d_2$ logarithmically and thus satisfies our requirement. Therefore, our bound correctly delineate the order of the regret.

- **Group-sparse Matrix Bandit.** We first randomly generate s row indices from the set $[d_2]$, where $d_2 = 5$ and then we generate our true Θ^* by setting the selected rows in Θ^* to be non-zero and generated uniform randomly from $(0, 1)$. The other rows are set to be zero. We then perform the normalization. We set $K = 10$ so that there are ten different contexts available at each round. The contexts $\{X_{t,a_i}\}_{i=1}^K$ are generated from the zero-mean and identity-covariance normal distribution. We choose three different row dimensions $d_1 = 5, 10, 100$. The group-sparsity is chosen to be $s = 3$ in all three settings. The experiment results are shown in Figure 3. As we can observe, the figures show that the regret is at most a constant times $d_2^{1/3}s^{1/3}T^{2/3}\sqrt{\log(d_1T)}$, which only depends on the dimension d_1 logarithmically and thus satisfies our requirement. Therefore, our bound correctly delineate the order of the regret.

In conclusions, all the experimental results validate all our claims in Section 4.

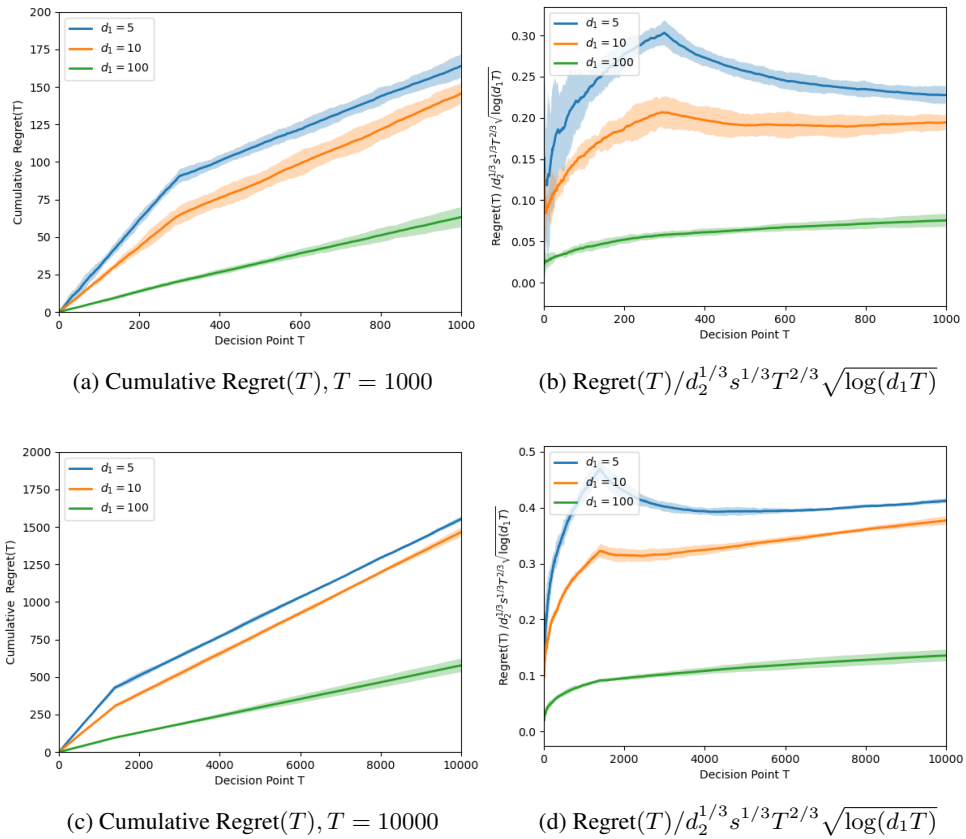


Figure 3. The experimental results in Group-sparse Matrix Bandit problem