
Multi-task Representation Learning for Pure Exploration in Bilinear Bandits

Subhojyoti Mukherjee
ECE Department
UW-Madison
Wisconsin, Madison
smukherjee27@wisc.edu

Qiaomin Xie
ISyE Department
UW-Madison
Wisconsin, Madison

Josiah P. Hanna
CS Department
UW-Madison
Wisconsin, Madison

Robert Nowak
ECE Department
UW-Madison
Wisconsin, Madison

Abstract

We study multi-task representation learning for the problem of pure exploration in bilinear bandits. In bilinear bandits, an action takes the form of a pair of arms from two different entity types and the reward is a bilinear function of the known feature vectors of the arms. In the *multi-task bilinear bandit problem*, we aim to find optimal actions for multiple tasks that share a common low-dimensional linear representation. The objective is to leverage this characteristic to expedite the process of identifying the best pair of arms for all tasks. We propose the algorithm **GOBLIN** that uses an experimental design approach to optimize sample allocations for learning the global representation as well as minimize the number of samples needed to identify the optimal pair of arms in individual tasks. To the best of our knowledge, this is the first study to give sample complexity analysis for pure exploration in bilinear bandits with shared representation. Our results demonstrate that by learning the shared representation across tasks, we achieve significantly improved sample complexity compared to the traditional approach of solving tasks independently.

1 Introduction

Bilinear bandits (Jun et al., 2019; Lu et al., 2021; Kang et al., 2022) are an important class of sequential decision-making problems. In bilinear bandits (as opposed to the standard linear bandit setting) we are given a pair of arms $\mathbf{x}_t \in \mathbb{R}^{d_1}$ and $\mathbf{z}_t \in \mathbb{R}^{d_2}$ at every round t and the interaction of this pair of arms with a low-rank hidden parameter, $\Theta_* \in \mathbb{R}^{d_1 \times d_2}$ generates the noisy feedback (reward) $r_t = \mathbf{x}_t^\top \Theta_* \mathbf{z}_t + \eta_t$. The η_t is random 1-subGaussian noise.

A lot of real-world applications exhibit the above bilinear feedback structure, particularly applications that involve selecting pairs of items and evaluating their compatibility. For example, in a drug discovery application, scientists may want to determine whether a particular (drug, protein) pair interacts in the desired way (Luo et al., 2017). Likewise, an online dating service might match a pair of people and gather feedback about their compatibility (Shen et al., 2023). A clothing website’s recommendation system may suggest a pair of items (top, bottom) for a customer based on their likelihood of matching (Reyes et al., 2021). In all of these scenarios, the two items are considered as a single unit, and the system must utilize available feature vectors $(\mathbf{x}_t, \mathbf{z}_t)$ to learn which features of the pairs are most indicative of positive feedback in order to make effective recommendations. All the previous works in this setting (Jun et al., 2019; Lu et al., 2021; Kang et al., 2022) exclusively

focused on maximizing the number of pairs with desired interactions discovered over time (regret minimization). However, in many real-world applications where obtaining a sample is expensive and time-consuming, e.g., clinical trials (Zhao et al., 2009; Zhang et al., 2012), it is often desirable to identify the optimal option using as few samples as possible, i.e., we face the pure exploration scenario (Fiez et al., 2019; Katz-Samuels et al., 2020) rather than regret minimization.

Moreover, in various decision-making scenarios, we may encounter multiple interrelated tasks such as treatment planning for different diseases (Bragman et al., 2018) and content optimization for multiple websites (Agarwal et al., 2009). Often, there exists a shared representation among these tasks, such as the features of drugs or the representations of website items. Therefore, we can leverage this shared representation to accelerate learning. This area of research is called multi-task representation learning and has recently generated a lot of attention in machine learning (Bengio et al., 2013; Li et al., 2014; Maurer et al., 2016; Du et al., 2020; Tripuraneni et al., 2021). There are many applications of this multi-task representation learning in real-world settings. For instance, in clinical treatment planning, we seek to determine the optimal treatments for multiple diseases, and there may exist a low-dimensional representation common to multiple diseases. To avoid the time-consuming process of conducting clinical trials for individual tasks and collecting samples, we utilize the shared representation and decrease the number of required samples.

The above multi-task representation learning naturally shows up in bilinear bandit setting as follows: Let there be M tasks indexed as $m = 1, 2, \dots, M$ with each task having its own hidden parameter $\Theta_{m,*} \in \mathbb{R}^{d_1 \times d_2}$. Let each $\Theta_{m,*}$ has a decomposition of $\Theta_{m,*} = \mathbf{B}_1 \mathbf{S}_{m,*} \mathbf{B}_2^\top$, where $\mathbf{B}_1 \in \mathbb{R}^{d_1 \times k_1}$ and $\mathbf{B}_2 \in \mathbb{R}^{d_2 \times k_2}$ are shared across tasks, but $\mathbf{S}_{m,*} \in \mathbb{R}^{k_1 \times k_2}$ is specific for task m . We assume that $k_1, k_2 \ll d_1, d_2$ and $M \gg d_1, d_2$. Thus, \mathbf{B}_1 and \mathbf{B}_2 provide a means of dimensionality reduction. Furthermore, we assume that each $\mathbf{S}_{m,*}$ has rank $r \ll \min\{k_1, k_2\}$. In the terminology of multi-task representation learning $\mathbf{B}_1, \mathbf{B}_2$ are called *feature extractors* and $\mathbf{x}_{m,t}, \mathbf{z}_{m,t}$ are called *rich observations* (Yang et al., 2020, 2022; Du et al., 2023). The reward for the task $m \in \{1, 2, \dots, M\}$ at round t is

$$r_{m,t} = \mathbf{x}_{m,t}^\top \Theta_{m,*} \mathbf{z}_{m,t} + \eta_{m,t} = \underbrace{\mathbf{x}_{m,t}^\top \mathbf{B}_1}_{\mathbf{g}_{m,t}^\top} \mathbf{S}_{m,*} \underbrace{\mathbf{B}_2^\top \mathbf{z}_{m,t}}_{\mathbf{v}_{m,t}} + \eta_{m,t} = \mathbf{g}_{m,t}^\top \mathbf{S}_{m,*} \mathbf{v}_{m,t} + \eta_{m,t}. \quad (1)$$

Observe that similar to the learning procedure in Yang et al. (2020, 2022), at each round $t = 1, 2, \dots$, for each task $m \in [M]$, the learner selects a left and right action $\mathbf{x}_{m,t} \in \mathcal{X}$ and $\mathbf{z}_{m,t} \in \mathcal{Z}$. After the player commits the batch of actions for each task $\{\mathbf{x}_{m,t}, \mathbf{z}_{m,t} : m \in [M]\}$, it receives the batch of rewards $\{r_{m,t} : m \in [M]\}$. Also note that in (1) we define the $\tilde{\mathbf{g}}_{m,t} \in \mathbb{R}^{k_1}, \tilde{\mathbf{v}}_{m,t} \in \mathbb{R}^{k_2}$ as the latent features, and both $\tilde{\mathbf{g}}_{m,t}, \tilde{\mathbf{v}}_{m,t}$ are unknown to the learner and needs to be learned for each task m (hence the name multi-task representation learning).

In this paper, we focus on pure exploration for multi-task representation learning in bilinear bandits where the goal is to find the optimal left arm $\mathbf{x}_{m,*}$ and right arm $\mathbf{z}_{m,*}$ for each task m with a minimum number of samples (fixed confidence setting). First, consider a single-task setting and let Θ_* have low rank r . Let the SVD of the $\Theta_* = \mathbf{U} \mathbf{D} \mathbf{V}^\top$. Prima-facie, if \mathbf{U} and \mathbf{V} are known then one might want to project all the left and right arms in the $r \times r$ subspace of \mathbf{U} and \mathbf{V} and reduce the bilinear bandit problem into a r^2 dimension linear bandit setting. Then one can apply one of the algorithms from Soare et al. (2014); Fiez et al. (2019); Katz-Samuels et al. (2020) to solve this r^2 dimensional linear bandit pure exploration problem. Following the analysis of this line of work (in linear bandits) (Mason et al., 2021; Mukherjee et al., 2022, 2023) one might conjecture that a sample complexity bound of $\tilde{O}(r^2/\Delta^2)$ is possible where Δ is the minimum reward gap and $\tilde{O}(\cdot)$ hides log factors. Similarly, for the multi-task setting one might be tempted to use the linear bandit analysis of Du et al. (2023) to convert this problem into M concurrent r^2 dimensional linear bandit problems with shared representation and achieve a sample complexity bound of $\tilde{O}(Mr^2/\Delta^2)$. However, these matrices (subspaces) are unknown and so there is a model mismatch as noted in the regret analysis of bilinear bandits (Jun et al., 2019; Lu et al., 2021; Kang et al., 2022). Thus it is difficult to apply the r^2 dimensional linear bandit sample complexity analysis. Following the regret analysis of bilinear bandit setting by Jun et al. (2019); Lu et al. (2021); Kang et al. (2022) we know that the effective dimension is actually $(d_1 + d_2)r$. Similarly for the multi-task representation learning the effective dimension should scale with the learned latent features $(k_1 + k_2)r$. Hence the natural questions to ask are these:

- 1) Can we design a single-task pure exploration bilinear bandit algorithm whose sample complexity scales as $\tilde{O}((d_1 + d_2)r/\Delta^2)$?

2) Can we design an algorithm for multi-task pure exploration bilinear bandit problem that can learn the latent features and has sample complexity that scales as $\tilde{O}(M(k_1 + k_2)r/\Delta^2)$?

In this paper, we answer both these questions affirmatively. In doing so, we make the following novel contributions to the growing literature of multi-task representation learning in online settings:

- 1) We formulate the multi-task bilinear representation learning problem. To our knowledge, this is the first work that explores pure exploration in a multi-task bilinear representation learning setting.
- 2) We proposed the algorithm **GOBLIN** for a single-task pure exploration bilinear bandit setting whose sample complexity scales as $\tilde{O}((d_1 + d_2)r/\Delta^2)$. This improves over RAGE (Fiez et al., 2019) whose sample complexity scales as $\tilde{O}((d_1 d_2)/\Delta^2)$.
- 3) Our algorithm **GOBLIN** for multi-task pure exploration bilinear bandit problem learns the latent features and has sample complexity that scales as $\tilde{O}(M(k_1 + k_2)r/\Delta^2)$. This improves over DouExpDes (Du et al., 2023) whose samples complexity scales as $\tilde{O}(M(k_1 k_2)/\Delta^2)$.

Preliminaries: We assume that $\|\mathbf{x}\|_2 \leq 1$, $\|\mathbf{z}\|_2 \leq 1$, $\|\Theta_*\|_F \leq S_0$ and the r -th largest singular value of $\Theta_* \in \mathbb{R}^{d_1 \times d_2}$ is S_r . Let $p := d_1 d_2$ denote the ambient dimension, and $k = (d_1 + d_2)r$ denote the effective dimension. Let $[n] := \{1, 2, \dots, n\}$. Let $\mathbf{x}_*, \mathbf{z}_* := \arg \max_{\mathbf{x}, \mathbf{z}} \mathbf{x}^\top \Theta_* \mathbf{z}$. For any \mathbf{x}, \mathbf{z} define the gap $\Delta(\mathbf{x}, \mathbf{z}) := \mathbf{x}_*^\top \Theta_* \mathbf{z}_* - \mathbf{x}^\top \Theta_* \mathbf{z}$ and furthermore $\Delta = \min_{\mathbf{x} \neq \mathbf{x}_*, \mathbf{z} \neq \mathbf{z}_*} \Delta(\mathbf{x}, \mathbf{z})$. Similarly, for any arbitrary vector $\mathbf{w} \in \mathcal{W}$ define the gap of $\mathbf{w} \in \mathbb{R}^p$ as $\Delta(\mathbf{w}) := (\mathbf{w}_* - \mathbf{w})^\top \theta_*$, for some $\theta_* \in \mathbb{R}^p$ and furthermore, $\Delta = \min_{\mathbf{w} \neq \mathbf{w}_*} \Delta(\mathbf{w})$. If $\mathbf{A} \in \mathbb{R}_{\geq 0}^{d \times d}$ is a positive semidefinite matrix, and $\mathbf{w} \in \mathbb{R}^p$ is a vector, let $\|\mathbf{w}\|_{\mathbf{A}}^2 := \mathbf{w}^\top \mathbf{A} \mathbf{w}$ denote the induced semi-norm. Given any vector $\mathbf{b} \in \mathbb{R}^{|\mathcal{W}|}$ we denote the \mathbf{w} -th component as $\mathbf{b}_{\mathbf{w}}$. Let $\Delta_{\mathcal{W}} := \{\mathbf{b} \in \mathbb{R}^{|\mathcal{W}|} : \mathbf{b}_{\mathbf{w}} \geq 0, \sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} = 1\}$ denote the set of probability distributions on \mathcal{W} . We define $\mathcal{Y}(\mathcal{W}) = \{\mathbf{w} - \mathbf{w}' : \forall \mathbf{w}, \mathbf{w}' \in \mathcal{W}, \mathbf{w} \neq \mathbf{w}'\}$ as the directions obtained from the differences between each pair of arms and $\mathcal{Y}^*(\mathcal{W}) = \{\mathbf{w}_* - \mathbf{w} : \forall \mathbf{w} \in \mathcal{W} \setminus \mathbf{w}_*\}$ as the directions obtained from the differences between the optimal arm and each suboptimal arm.

2 Pure Exploration in Single-Task Bilinear Bandits

In this section, we consider pure exploration in a single-task bilinear bandit setting as a warm-up to the main goal of learning representations for the multi-task bilinear bandit. To our knowledge, this is the first study of pure exploration in single-task bilinear bandits. We first recall the single-task bilinear bandit setting as follows: At every round $t = 1, 2, \dots$ the learner observes the reward $r_t = \mathbf{x}_t^\top \Theta_* \mathbf{z}_t + \eta_t$ where the low rank hidden parameter $\Theta_* \in \mathbb{R}^{d_1 \times d_2}$ is unknown to the learner, $\mathbf{x}_t \in \mathbb{R}^{d_1}$, $\mathbf{z}_t \in \mathbb{R}^{d_2}$ are visible to the learner, and η_t is a 1-sub-Gaussian noise. We assume that the matrix Θ_* has a low rank r which is known to the learner and $d_1, d_2 \gg r$. Finally recall that the goal is to identify the optimal left and right arms $\mathbf{x}_*, \mathbf{z}_*$ with a minimum number of samples.

We propose a phase-based, two-stage arm elimination algorithm called **G-Optimal Design for Bilinear Bandits** (abbreviated as **GOBLIN**). **GOBLIN** proceeds in phases indexed by $\ell = 1, 2, \dots$. As this is a pure-exploration problem, the total number of samples is controlled by the total phases which depends on the intrinsic problem complexity. Each phase ℓ of **GOBLIN** consists of two stages; the estimation of Θ_* stage, which runs for τ_ℓ^E rounds, and pure exploration in rotated arms stage that runs for τ_ℓ^G rounds. We will define τ_ℓ^E in Section 2.1, while rotated arms and τ_ℓ^G are defined in Section 2.2. At the end of every phase, **GOBLIN** eliminates sub-optimal arms to build the active set for the next phase and stops when only the optimal left and right arms are remaining. Now we discuss the individual stages that occur at every phase ℓ of **GOBLIN**.

2.1 Estimating Subspaces of Θ_* (Stage 1 of the ℓ -th phase)

In the first stage of phase ℓ , **GOBLIN** estimates the row and column sub-spaces Θ_* . Then **GOBLIN** uses these estimates to reduce the bilinear bandit problem in the original ambient dimension $p := d_1 d_2$ to a lower effective dimension $k := (d_1 + d_2)r$. To do this, **GOBLIN** first vectorizes the $\mathbf{x} \in \mathbb{R}^{d_1}, \mathbf{z} \in \mathbb{R}^{d_2}$ into a new vector $\bar{\mathbf{w}} \in \mathbb{R}^p$ and then solves the E -optimal design in Step 3 of

Algorithm 1 (Pukelsheim, 2006; Jun et al., 2019; Du et al., 2023). Let the solution to the E -optimal design problem at the stage 1 of ℓ -th phase be denoted by \mathbf{b}_ℓ^E . Then **GOBLIN** samples each $\bar{\mathbf{w}}$ for $\lceil \tau_\ell^E \mathbf{b}_\ell^E \rceil$ times, where $\tau_\ell^E = \tilde{O}(\sqrt{d_1 d_2 r} / S_r)$ (step 7 of Algorithm 1). In this paper, we sample an arm $\lceil \tau_\ell^E \mathbf{b}_\ell^E \rceil$ number of times. However, this may lead to over-sampling of an arm than what the design (G or E -optimal) is actually suggesting. However, we can match the number of allocations of an arm to the design using an *efficient Rounding Procedures* (see Pukelsheim (2006); Fiez et al. (2019)). Let $\hat{\Theta}_\ell$ be estimate of Θ_* in stage 1 of phase ℓ . **GOBLIN** estimates this by solving the following well-defined regularized minimization problem with nuclear norm penalty:

$$\hat{\Theta}_\ell = \arg \min_{\Theta \in \mathbb{R}^{d_1 \times d_2}} L_\ell(\Theta) + \gamma_\ell \|\Theta\|_{\text{nuc}}, \quad L_\ell(\Theta) = \langle \Theta, \Theta \rangle - \frac{2}{\tau_\ell^E} \sum_{s=1}^{\tau_\ell^E} \langle \tilde{\psi}_\nu(r_s \cdot Q(\mathbf{x}_s \mathbf{z}_s^\top)), \Theta \rangle \quad (2)$$

where $Q(\cdot)$, $\tilde{\psi}_\nu(\cdot)$, are appropriate functions stated in Definition 1, 3 respectively in Appendix A.3. The $Q(\cdot)$ function takes as input the rank-one matrix $\mathbf{x}_s \mathbf{z}_s^\top$ which is obtained after reshaping $\bar{\mathbf{w}}_s$. Note that \mathbf{x}_s , and \mathbf{z}_s are the observed vectors in d_1 and d_2 dimension and $\hat{\Theta}_\ell \in \mathbb{R}^{d_1 \times d_2}$. Finally, set the regularization parameter $\gamma_\ell := 4\sqrt{\frac{2(4+S_0^2)C d_1 d_2 \log(2(d_1+d_2)/\delta)}{\tau_\ell^E}}$. This is in step 8 of Algorithm 1.

2.2 Optimal Design for Rotated Arms (Stage 2 of ℓ -th phase)

In stage 2 of phase ℓ , **GOBLIN** leverages the information about the learned sub-space of Θ_* to rotate the arm set and then run the optimal design on the rotated arm set. Once we recover $\hat{\Theta}_\ell$, one might be tempted to run a pure exploration algorithm (Soare et al., 2014; Fiez et al., 2019; Katz-Samuels et al., 2020; Zhu et al., 2021) to identify \mathbf{x}_* and \mathbf{z}_* . However, then the sample complexity will scale with $d_1 d_2$. In contrast **GOBLIN** uses the information about the learned sub-space of Θ_* to reduce the problem from ambient dimension $d_1 d_2$ to effective dimension $(d_1 + d_2)r$. This reduction is done as follows: Let $\hat{\Theta}_\ell = \hat{\mathbf{U}}_\ell \hat{\mathbf{D}}_\ell \hat{\mathbf{V}}_\ell^\top$ be the SVD of $\hat{\Theta}_\ell$ in the ℓ -th phase. Let $\hat{\mathbf{U}}_\ell^\perp$ and $\hat{\mathbf{V}}_\ell^\perp$ be orthonormal bases of the complementary subspaces of $\hat{\mathbf{U}}_\ell$ and $\hat{\mathbf{V}}_\ell$ respectively. Let \mathcal{X}_ℓ and \mathcal{Z}_ℓ be the active set of arms in the stage 2 of phase ℓ . Then rotate the arm sets such that new rotated arm sets are as follows:

$$\mathcal{X}_\ell = \{\mathbf{x} = [\hat{\mathbf{U}}_\ell \hat{\mathbf{U}}_\ell^\perp]^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X}_\ell\}, \mathcal{Z}_\ell = \{\mathbf{z} = [\hat{\mathbf{V}}_\ell \hat{\mathbf{V}}_\ell^\perp]^\top \mathbf{z} \mid \mathbf{z} \in \mathcal{Z}_\ell\}. \quad (3)$$

Let $\hat{\mathbf{H}}_\ell = [\hat{\mathbf{U}}_\ell \hat{\mathbf{U}}_\ell^\perp]^\top \hat{\Theta}_\ell [\hat{\mathbf{V}}_\ell \hat{\mathbf{V}}_\ell^\perp]$. Then define vectorized arm set so that the last $(d_1 - r) \cdot (d_2 - r)$ components are from the complementary subspaces as follows:

$$\begin{aligned} \mathcal{W}_\ell &= \{[\text{vec}(\mathbf{x}_{1:r} \mathbf{z}_{1:r}^\top); \text{vec}(\mathbf{x}_{r+1:d_1} \mathbf{z}_{1:r}^\top); \text{vec}(\mathbf{x}_{1:r} \mathbf{z}_{r+1:d_2}^\top); \\ &\quad \text{vec}(\mathbf{x}_{r+1:d_1} \mathbf{z}_{r+1:d_2}^\top)] \in \mathbb{R}^{d_1 d_2} : \mathbf{x} \in \mathcal{X}_\ell, \mathbf{z} \in \mathcal{Z}_\ell\} \\ \hat{\boldsymbol{\theta}}_{\ell,1:k} &= [\text{vec}(\hat{\mathbf{H}}_{\ell,1:r,1:r}); \text{vec}(\hat{\mathbf{H}}_{\ell,r+1:d_1,1:r}); \text{vec}(\hat{\mathbf{H}}_{\ell,1:r,r+1:d_2})], \\ \hat{\boldsymbol{\theta}}_{\ell,k+1:p} &= \text{vec}(\hat{\mathbf{H}}_{\ell,r+1:d_1,r+1:d_2}). \end{aligned} \quad (4)$$

which implies $\|\hat{\boldsymbol{\theta}}_{k+1:p}\|_2 = O(d_1 d_2 r / \tau_\ell^E)$ by Lemma 3 in Appendix A.1. So the last $p - k$ components of $\hat{\boldsymbol{\theta}}_\ell$ are very small compared to the first k components. Hence, **GOBLIN** has now reduced the $d_1 d_2$ dimensional linear bandit to $(d_1 + d_2)r$ dimensional linear bandit using (3), (4). This is shown in step 10 of Algorithm 1.

Now in stage 2 of phase ℓ , **GOBLIN** implements G -optimal design (Pukelsheim, 2006; Fiez et al., 2019) in the rotated arm set $\mathcal{X}_\ell, \mathcal{Z}_\ell$ defined in (3). To do this, first **GOBLIN** defines the rotated vector $\mathbf{w} = [\mathbf{x}_{1:d_1}; \mathbf{z}_{1:d_2}] \in \mathbb{R}^p$ that belong to the set \mathcal{W}_ℓ . Then **GOBLIN** solves the G -optimal design (Pukelsheim, 2006) as follows:

$$\hat{\mathbf{b}}_\ell^G = \arg \min_{\mathbf{b}_\mathbf{w}} \max_{\mathbf{w}, \mathbf{w}' \in \mathcal{W}_\ell} \|\mathbf{w} - \mathbf{w}'\|^2_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_\mathbf{w} \mathbf{w} \mathbf{w}^\top + \Lambda_\ell / n)^{-1}}. \quad (5)$$

This is shown in step 11 of Algorithm 1 and Λ_ℓ is defined in (6). It can be shown that sampling according to $\hat{\mathbf{b}}_\ell^G$ leads to the optimal sample complexity. This is discussed in Remark 1 in Appendix A.2. The key point to note from (5) is that due to the estimation in the rotated arm space \mathcal{W}_ℓ we are guaranteed that the support of $\text{supp}(\hat{\mathbf{b}}_\ell^G) \leq \tilde{O}(k(k+1)/2)$ (Pukelsheim, 2006). On the other hand,

if the G-optimal design of Fiez et al. (2019); Katz-Samuels et al. (2020) are run in $d_1 d_2$ dimension then the support of $\hat{\mathbf{b}}_\ell^G$ will scale with $d_1 d_2$ which will lead to higher sample complexity. Then **GOBLIN** samples each $\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell$ for $\lceil \tau_\ell^G \mathbf{b}_{\ell, \underline{\mathbf{w}}}^G \rceil$ times, where $\tau_\ell^G := \lceil \frac{8B_\ell^* \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \log(4\ell^2 |\mathcal{W}|/\delta)}{\epsilon_\ell^2} \rceil$. Note that the total length of phase ℓ , combining stages 1 and 2 is $(\tau_\ell^E + \tau_\ell^G)$ rounds. Observe that the stage 1 design is on the whole arm set $\overline{\mathcal{W}}$ whereas stage 2 design is on the refined active set $\underline{\mathcal{W}}_\ell$.

Let the observed features in stage 2 of phase ℓ be denoted by $\underline{\mathbf{W}}_\ell \in \mathbb{R}^{\tau_\ell^G \times p}$, and $\mathbf{r}_\ell \in \mathbb{R}^{\tau_\ell^G}$ be the observed rewards. Define the diagonal matrix $\mathbf{\Lambda}_\ell$ as

$$\mathbf{\Lambda}_\ell = \text{diag}[\underbrace{\lambda, \dots, \lambda}_k, \underbrace{\lambda_\ell^\perp, \dots, \lambda_\ell^\perp}_{p-k}] \quad (6)$$

where, $\lambda_\ell^\perp := \tau_{\ell-1}^G / 8k \log(1 + \tau_{\ell-1}^G / \lambda) \gg \lambda$. Deviating from Soare et al. (2014); Fiez et al. (2019) **GOBLIN** constructs a regularized least square estimator at phase ℓ as follows

$$\hat{\boldsymbol{\theta}}_\ell = \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^p} \frac{1}{2} \|\underline{\mathbf{W}}_\ell \boldsymbol{\theta} - \mathbf{r}_\ell\|_2^2 + \frac{1}{2} \|\boldsymbol{\theta}\|_{\mathbf{\Lambda}_\ell}^2. \quad (7)$$

This regularized least square estimator in (7) forces the last $p - k$ components of $\hat{\boldsymbol{\theta}}_\ell$ to be very small compared to the first k components. Then **GOBLIN** builds the estimate $\hat{\boldsymbol{\theta}}_\ell$ from (7) only from the observations from this phase (step 13 in Algorithm 1) and eliminates sub-optimal actions in step 14 in Algorithm 1 using the estimator $\hat{\boldsymbol{\theta}}_\ell$. Finally **GOBLIN** eliminates sub-optimal arms to build the next phase active set $\underline{\mathcal{W}}_\ell$ and stops when $|\underline{\mathcal{W}}_\ell| = 1$. **GOBLIN** outputs the arm in $\underline{\mathcal{W}}_\ell$ and reshapes it to get the $\hat{\mathbf{x}}_*$ and $\hat{\mathbf{z}}_*$. The full pseudocode is presented in Algorithm 1.

Algorithm 1 G-Optimal Design for Bilinear Bandits (**GOBLIN**) for single-task setting

- 1: Input: arm set \mathcal{X}, \mathcal{Z} , confidence δ , rank r of $\boldsymbol{\Theta}_*$, spectral bound S_r of $\boldsymbol{\Theta}_*$, $S, S_\ell^\perp := \frac{8d_1 d_2 r}{\tau_\ell^E S_r^2} \log\left(\frac{d_1 + d_2}{\delta_\ell}\right)$, $\lambda, \lambda_\ell^\perp := \tau_{\ell-1}^G / 8(d_1 + d_2)r \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})$. Let $p := d_1 d_2$, $k := (d_1 + d_2)r$.
 - 2: Let $\underline{\mathcal{W}}_1 \leftarrow \underline{\mathcal{W}}, \ell \leftarrow 1, \tau_0^G := \log(4\ell^2 |\mathcal{X}|/\delta)$. Define $\mathbf{\Lambda}_\ell$ as in (6), $B_\ell^* := (8\sqrt{\lambda} S + \sqrt{\lambda_\ell^\perp S_\ell^\perp})$.
 - 3: Define a vectorized arm $\overline{\mathbf{w}} := [\mathbf{x}_{1:d_1}; \mathbf{z}_{1:d_2}]$ and $\overline{\mathbf{w}} \in \overline{\mathcal{W}}$. Let $\tau_\ell^E := \frac{\sqrt{8d_1 d_2 r \log(4\ell^2 |\mathcal{W}|/\delta_\ell)}}{S_r}$. Let the E -optimal design be $\mathbf{b}_\ell^E := \arg \min_{\mathbf{b} \in \Delta_{\overline{\mathcal{W}}}} \|(\sum_{\overline{\mathbf{w}} \in \overline{\mathcal{W}}} \mathbf{b}_{\overline{\mathbf{w}}} \overline{\mathbf{w}} \overline{\mathbf{w}}^\top)^{-1}\|$.
 - 4: **while** $|\underline{\mathcal{W}}_\ell| > 1$ **do**
 - 5: $\epsilon_\ell = 2^{-\ell}, \delta_\ell = \delta/\ell^2$.
 - 6: **(Stage 1:) Explore the Low-Rank Subspace**
 - 7: Pull arm $\overline{\mathbf{w}} \in \overline{\mathcal{W}}$ exactly $\lceil \hat{\mathbf{b}}_{\ell, \overline{\mathbf{w}}}^E \tau_\ell^E \rceil$ times and observe rewards r_t , for $t = 1, \dots, \tau_\ell^E$.
 - 8: Compute $\hat{\boldsymbol{\theta}}_\ell$ using (2).
 - 9: **(Stage 2:) Reduction to low dimensional linear bandits**
 - 10: Let the SVD of $\hat{\boldsymbol{\theta}}_\ell = \hat{\mathbf{U}}_\ell \hat{\mathbf{D}}_\ell \hat{\mathbf{V}}_\ell^\top$. Rotate arms in active set $\underline{\mathcal{W}}_{\ell-1}$ to build $\underline{\mathcal{W}}_\ell$ following (4).
 - 11: Let $\hat{\mathbf{b}}_\ell^G := \arg \min_{\mathbf{b}_{\underline{\mathbf{w}}}} \max_{\mathbf{w}, \mathbf{w}' \in \underline{\mathcal{W}}_\ell} \|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|_{(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell/n)^{-1}}$.
 - 12: Define $\rho^G(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) := \min_{\mathbf{b}_{\underline{\mathbf{w}}}} \max_{\mathbf{w}, \mathbf{w}' \in \underline{\mathcal{W}}_\ell} \|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|_{(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell/n)^{-1}}$.
 - 13: Set $\tau_\ell^G := \lceil \frac{64B_\ell^* \rho^G(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) \log(4\ell^2 |\mathcal{W}|/\delta_\ell)}{\epsilon_\ell^2} \rceil$. Then pull arm $\underline{\mathbf{w}} \in \underline{\mathcal{W}}$ exactly $\lceil \hat{\mathbf{b}}_{\ell, \underline{\mathbf{w}}}^G \tau_\ell^G \rceil$ times and construct the least squares estimator $\hat{\boldsymbol{\theta}}_\ell$ using only the observations of this phase where $\hat{\boldsymbol{\theta}}_\ell$ is defined in (7). Note that $\hat{\boldsymbol{\theta}}_\ell$ is also rotated following (4).
 - 14: Eliminate arms such that $\underline{\mathcal{W}}_{\ell+1} \leftarrow \underline{\mathcal{W}}_\ell \setminus \{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell : \max_{\mathbf{w}' \in \underline{\mathcal{W}}_\ell} \langle \underline{\mathbf{w}}' - \underline{\mathbf{w}}, \hat{\boldsymbol{\theta}}_\ell \rangle > 2\epsilon_\ell\}$
 - 15: $\ell \leftarrow \ell + 1$
 - 16: Output the arm in $\underline{\mathcal{W}}_\ell$ and reshape to get the $\hat{\mathbf{x}}_*$ and $\hat{\mathbf{z}}_*$
-

2.3 Sample Complexity Analysis of Single-Task **GOBLIN**

We now analyze the sample complexity of **GOBLIN** in the single-task setting through the following theorem.

Theorem 1. (informal) With probability at least $1 - \delta$, **GOBLIN** returns the best arms \mathbf{x}_* , \mathbf{z}_* , and the number of samples used is bounded by $\tilde{O}\left(\frac{(d_1+d_2)r}{\Delta^2} + \frac{\sqrt{d_1 d_2 r}}{S_r}\right)$.

Discussion 1. In Theorem 1 the first quantity is the number of samples needed to identify the best arms \mathbf{x}_* , \mathbf{z}_* while the second quantity is the number of samples to learn Θ_* (which is required to find the best arms). Note that the magnitude of S_r would be free of d_1, d_2 since Θ_* contains only r nonzero singular values and $\|\Theta_*\| \leq 1$, and hence we assume that $S_r = \Theta(1/\sqrt{r})$ (Kang et al., 2022). So the sample complexity of single-task **GOBLIN** scales as $\tilde{O}(\frac{(d_1+d_2)r}{\Delta^2})$. However, if one runs RAGE (Fiez et al., 2019) on the arms in \mathcal{X}, \mathcal{Z} then the sample complexity will scale as $\tilde{O}(\frac{d_1 d_2}{\Delta^2})$.

Proof (Overview) of Theorem 1: Step 1 (Subspace estimation in high dimension): We denote the vectorized arms in high dimension as $\bar{\mathbf{w}} \in \bar{\mathcal{W}}$. We run the E -optimal design to sample the arms in $\bar{\mathcal{W}}$. Note that this E -optimal design satisfies the distribution assumption of Kang et al. (2022) which enables us to apply the Lemma 3 in Appendix A.1. This leads to $\|\hat{\Theta}_\ell - \Theta_*\|_F^2 \leq \frac{C_1 d_1 d_2 r \log(2(d_1+d_2)/\delta)}{\tau_\ell^E}$ for some $C_1 > 0$. Also, note that in the first stage of the ℓ -th phase by setting $\tau_\ell^E = \frac{\sqrt{8d_1 d_2 r \log(4\ell^2 |\mathcal{W}|/\delta_\ell)}}{S_r}$ and sampling each arm $\bar{\mathbf{w}} \in \bar{\mathcal{W}}$ exactly $\lceil \hat{\mathbf{b}}_{\ell, \bar{\mathbf{w}}}^E \tau_\ell^E \rceil$ times we are guaranteed that $\|\theta_{k+1:p}^*\|_2 = O(d_1 d_2 r / \tau_\ell^E)$. Summing up over $\ell = 1$ to $\lceil \log_2(4\Delta^{-1}) \rceil$ we get that the total sample complexity of the first stage is bounded by $\tilde{O}(\sqrt{d_1 d_2 r} / S_r)$.

Step 2 (Effective dimension for rotated arms): We rotate the arms $\bar{\mathbf{w}} \in \bar{\mathcal{W}}$ in high dimension to get the rotated arms $\mathbf{w} \in \mathcal{W}_\ell$ in step 10 of Algorithm 1. Then we show that the effective dimension of \mathbf{w} scales $8k \log(1 + \tau_{\ell-1}^G / \lambda)$ when $\lambda_\ell^\perp = \frac{\tau_{\ell-1}^G}{8k \log(1 + \tau_{\ell-1}^G / \lambda)}$ in Lemma 7 of Appendix A.4. Note that this requires a different proof technique than Valko et al. (2014) where the budget n is given apriori and effective dimension scales with $\log(n)$. This step also diverges from the pure exploration proof technique of Fiez et al. (2019); Katz-Samuels et al. (2020) as there is no parameter λ_ℓ^\perp to control during phase ℓ , and the effective dimensions in those papers do not depend on phase length.

Step 3 (Bounded Support): For any phase ℓ , we can show that $1 \leq \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \leq p/\gamma_Y^2$ where, $\gamma_Y = \max\{c > 0 : c\mathcal{Y} \subset \text{conv}(\mathcal{W} \cup -\mathcal{W})\}$ is the gauge norm of \mathcal{Y} (Rockafellar, 2015). Note that this is a worst-case dependence when $\rho^G(\mathcal{Y}(\mathcal{W}_\ell))$ scales with p . Substituting this value of $\rho^G(\mathcal{Y}(\mathcal{W}_\ell))$ in the definition of λ_ℓ^\perp we can show that Λ_ℓ does not depend on \mathbf{w} or $\mathbf{y} = \mathbf{w} - \mathbf{w}'$. Then following Theorem 21.1 in Lattimore and Szepesvári (2020) we can show that the G -optimal design $\hat{\mathbf{b}}_\ell^G$ is equivalent to D -optimal design $\hat{\mathbf{b}}_\ell^D = \arg \max_{\mathbf{b}} \log \frac{|\sum_{\mathbf{w} \in \mathcal{W}_\ell} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda_\ell|}{|\Lambda_\ell|}$. Then using Frank-Wolfe algorithm (Jamieson and Jain, 2022) we can show the support $\hat{\mathbf{b}}_\ell^G$ or equivalently $\hat{\mathbf{b}}_\ell^D$ is bounded by at most $\frac{8k \log(1 + \tau_{\ell-1}^G / \lambda)(8k \log(1 + \tau_{\ell-1}^G / \lambda) + 1)}{2}$. This is shown in Lemma 9 (Appendix A.4).

Step 4 (Phase length and Elimination): Using the Lemma 9, concentration Lemma 5, and using the log determinant inequality in Lemma 7 and Proposition 1 (Appendix A.4) we show that the phase length in the second stage is given by $\tau_\ell^G = \lceil \frac{8B_* \rho(\mathcal{Y}(\mathcal{W}_\ell)) \log(2|\mathcal{W}|/\delta)}{(\mathbf{x}^\top(\hat{\Theta}_\ell - \Theta_*)^2)} \rceil$. This is discussed in Discussion 3 (Appendix A.4). We show in Lemma 10 (Appendix A.4) that setting this phase length and sampling each active arm in \mathcal{W}_ℓ exactly $\lceil \hat{\mathbf{b}}_{\ell, \mathbf{w}}^G \tau_\ell^G \rceil$ times results in the elimination of sub-optimal actions with high probability.

Step 5 (Total Samples): We first show that the total samples in the second phase are bounded by $O(\frac{k}{\gamma_Y^2} \log(\frac{k \log_2(\Delta^{-1}) |\mathcal{W}|}{\delta}) \lceil \log_2(\Delta^{-1}) \rceil)$ where the effective dimension $k = (d_1 + d_2)r$. Finally, we combine the total samples of phase ℓ as $(\tau_\ell^E + \tau_\ell^G)$. The final sample complexity is given by summing over all phases from $\ell = 1$ to $\lceil \log_2(4\Delta^{-1}) \rceil$. The claim of the theorem follows by noting $\tilde{O}(k/\gamma_Y^2) \leq \tilde{O}(k/\Delta^2)$.

3 Multi-task Representation Learning

In this section, we extend **GOBLIN** to multi-task representation learning for the bilinear bandit setting. In the multi-task setting, we now have M tasks, where each task $m \in [M]$ has a reward model stated

in (1). The learning proceeds as follows: At each round $t = 1, 2, \dots$, for each task $m \in [M]$, the learner selects a left and right action $\mathbf{x}_{m,t} \in \mathcal{X}$ and $\mathbf{z}_{m,t} \in \mathcal{Z}$. After the player commits the batch of actions for each task $\{\mathbf{x}_{m,t}, \mathbf{z}_{m,t} : m \in [M]\}$, it receives the batch of rewards $\{r_{m,t} : m \in [M]\}$. Finally recall that the goal is to identify the optimal left and right arms $\mathbf{x}_{m,*}, \mathbf{z}_{m,*}$ for each task m with a minimum number of samples. We now state the following assumptions to enable representation learning across tasks.

Assumption 1. (Low-rank Tasks) We assume that the hidden parameter $\Theta_{m,*}$ for all the $m \in [M]$ have a decomposition $\Theta_{m,*} = \mathbf{B}_1 \mathbf{S}_{m,*} \mathbf{B}_2^\top$ and each $\mathbf{S}_{m,*}$ has rank r .

This is similar to the assumptions in Yang et al. (2020, 2022); Du et al. (2023) ensuring the feature extractors are shared across tasks in the bilinear bandit setting.

Assumption 2. (Diverse Tasks) We assume that $\sigma_{\min}(\frac{1}{M} \sum_{m=1}^M \Theta_{m,*}) \geq \frac{c_0}{S_r}$, for some $c_0 > 0$, S_r is the r -th largest singular value of $\Theta_{m,*}$ and $\sigma_{\min}(\mathbf{A})$ denotes the minimum eigenvalue of matrix \mathbf{A} .

This assumption is similar to the diverse tasks assumption of Yang et al. (2020, 2022); Tripuraneni et al. (2021); Du et al. (2023) and ensures the possibility of recovering the feature extractors \mathbf{B}_1 and \mathbf{B}_2 shared across tasks.

Our extension of GOBLIN to the multi-task setting is now a phase-based, *three-stage* arm elimination algorithm. In GOBLIN each phase $\ell = 1, 2, \dots$ consists of three stages; the stage for estimation of feature extractors $\mathbf{B}_1, \mathbf{B}_2$, which runs for τ_ℓ^E rounds, the stage for estimation of $\mathbf{S}_{m,*}$ which runs for $\sum_m \tilde{\tau}_{m,\ell}^E$ rounds, and a stage of pure exploration with rotated arms that runs for $\sum_m \tau_{m,\ell}^G$ rounds. We will define $\tau_{m,\ell}^E$ in Section 3.1, $\tilde{\tau}_{m,\ell}^E$ in Section 3.2, while the rotated arms and $\tau_{m,\ell}^G$ are defined in Section 3.3. At the end of every phase, GOBLIN eliminates sub-optimal arms to build the active set for the next phase and stops when only the optimal left and right arms are remaining. Now we discuss the individual stages that occur at every phase $\ell = 1, 2, \dots$ for multi-task GOBLIN.

3.1 Estimating Feature Extractors \mathbf{B}_1 and \mathbf{B}_2 (Stage 1 of Phase ℓ)

In the first stage of phase ℓ , GOBLIN leverages the batch of rewards $\{r_{m,t} : m \in [M]\}$ at every round t from M tasks to learn the feature extractors \mathbf{B}_1 and \mathbf{B}_2 . To do this, GOBLIN first vectorizes the $\mathbf{x} \in \mathcal{X}, \mathbf{z} \in \mathcal{Z}$ into a new vector $\bar{\mathbf{w}} = [\mathbf{x}_{1:d_1}; \mathbf{z}_{1:d_2}] \in \bar{\mathcal{W}}_m$ and then solves the E -optimal design in step 3 of Algorithm 2. Similar to the single-task setting (Section 2) GOBLIN samples each $\bar{\mathbf{w}} \in \bar{\mathcal{W}}_m$ for $\lceil \tau_\ell^E \mathbf{b}_{\ell,\bar{\mathbf{w}}}^E \rceil$ times for each task m , where $\tau_\ell^E = \tilde{O}(\sqrt{d_1 d_2 r} / S_r)$ and $\mathbf{b}_{\ell,\bar{\mathbf{w}}}^E$ is the solution to E -optimal design on $\bar{\mathbf{w}}$. Let the sampled arms for each task m at round s be denoted by $\mathbf{x}_{m,s}, \mathbf{z}_{m,s}$ which is obtained after reshaping $\bar{\mathbf{w}}_s$. Then it builds the estimator $\hat{\mathbf{Z}}_\ell$ as follows:

$$\begin{aligned} \hat{\mathbf{Z}}_\ell &= \arg \min_{\Theta \in \mathbb{R}^{d_1 \times d_2}} L_\ell(\Theta) + \gamma_\ell \|\Theta\|_{\text{nuc}}, \\ L_\ell(\Theta) &= \langle \Theta, \Theta \rangle - \frac{2}{M \tau_\ell^E} \sum_{m=1}^M \sum_{s=1}^{\tau_\ell^E} \langle \tilde{\psi}_\nu(r_{m,s} \cdot Q(\mathbf{x}_{m,s} \mathbf{z}_{m,s}^\top)), \Theta \rangle \end{aligned} \quad (8)$$

Then it performs SVD decomposition on $\hat{\mathbf{Z}}_\ell$, and let $\hat{\mathbf{B}}_1, \hat{\mathbf{B}}_2$ be the top- k_1 and top- k_2 left and right singular vectors of $\hat{\mathbf{Z}}_\ell$ respectively. These are the estimation of the feature extractors \mathbf{B}_1 and \mathbf{B}_2 .

3.2 Estimating Hidden Parameter $\mathbf{S}_{m,*}$ per Task (Stage 2 of phase ℓ)

In the second stage of phase ℓ , the goal is to recover the hidden parameter $\mathbf{S}_{m,*}$ for each task m . GOBLIN proceeds as follows: First, let $\tilde{\mathbf{g}}_m = \mathbf{x}^\top \hat{\mathbf{B}}_{1,\ell}$ and $\tilde{\mathbf{v}}_m = \mathbf{z}^\top \hat{\mathbf{B}}_{2,\ell}$ be the latent left and right arm respectively for each m . Then GOBLIN defines the vector $\tilde{\mathbf{w}} = [\tilde{\mathbf{g}}_m; \tilde{\mathbf{v}}_m] \in \tilde{\mathcal{W}}_m$ and then solves the E -optimal design in step 11 of Algorithm 2. It then samples for each task m , the latent arm $\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}_m$ for $\lceil \tilde{\tau}_{m,\ell}^E \tilde{\mathbf{b}}_{m,\ell,\tilde{\mathbf{w}}}^E \rceil$ times, where $\tilde{\tau}_{m,\ell}^E := \tilde{O}(\sqrt{k_1 k_2 r} / S_r)$ and $\tilde{\mathbf{b}}_{m,\ell,\tilde{\mathbf{w}}}^E$ is the solution to

E -optimal design on $\tilde{\mathbf{w}}$. Then it builds estimator $\hat{\mathbf{S}}_{m,\ell}$ for each task m in step 12 as follows:

$$\begin{aligned}\hat{\mathbf{S}}_{m,\ell} &= \arg \min_{\Theta \in \mathbb{R}^{k_1 \times k_2}} L'_\ell(\Theta) + \gamma_\ell \|\Theta\|_{\text{nuc}}, \\ L'_\ell(\Theta) &= \langle \Theta, \Theta \rangle - \frac{2}{\tilde{\tau}_{m,\ell}^E} \sum_{s=1}^{\tilde{\tau}_{m,\ell}^E} \langle \tilde{\psi}_\nu(r_{m,s} \cdot Q(\tilde{\mathbf{g}}_{m,s} \tilde{\mathbf{v}}_{m,s}^\top), \Theta) \rangle\end{aligned}\quad (9)$$

Once **GOBLIN** recovers the $\hat{\mathbf{S}}_{m,\ell}$ for each task m it has reduced the $d_1 d_2$ bilinear bandit to a $k_1 k_2$ dimension bilinear bandit where the left and right arms are $\tilde{\mathbf{g}}_m \in \mathcal{G}_m, \tilde{\mathbf{v}}_m \in \mathcal{V}_m$ respectively.

3.3 Optimal Design for Rotated Arms per Task (Stage 3 of phase ℓ)

In the third stage of phase ℓ , similar to Algorithm 1, the multi-task **GOBLIN** defines the rotated arm set $\underline{\mathcal{G}}_m, \underline{\mathcal{V}}_m$ for each task m for these $k_1 k_2$ dimensional bilinear bandits. Let the SVD of $\hat{\mathbf{S}}_{m,\ell} = \hat{\mathbf{U}}_{m,\ell} \hat{\mathbf{D}}_{m,\ell} \hat{\mathbf{V}}_{m,\ell}^\top$. Define $\hat{\mathbf{H}}_{m,\ell} = [\hat{\mathbf{U}}_{m,\ell} \hat{\mathbf{U}}_{m,\ell}^\perp]^\top \hat{\mathbf{S}}_{m,\ell} [\hat{\mathbf{V}}_{m,\ell} \hat{\mathbf{V}}_{m,\ell}^\perp]$. Then define the vectorized arm set so that the last $(k_1 - r) \cdot (k_2 - r)$ components are from the complementary subspaces as follows:

$$\begin{aligned}\underline{\mathcal{W}}_{m,\ell} &= \left\{ [\text{vec}(\tilde{\mathbf{g}}_{m,1:r} \tilde{\mathbf{v}}_{m,1:r}^\top); \text{vec}(\tilde{\mathbf{g}}_{m,r+1:k_1} \tilde{\mathbf{v}}_{m,1:r}^\top); \text{vec}(\tilde{\mathbf{g}}_{m,1:r} \tilde{\mathbf{v}}_{m,r+1:k_2}^\top); \right. \\ &\quad \left. \text{vec}(\tilde{\mathbf{g}}_{m,r+1:k_1} \tilde{\mathbf{v}}_{m,r+1:k_2}^\top)] \right\} \\ \hat{\boldsymbol{\theta}}_{m,\ell,1:k} &= [\text{vec}(\hat{\mathbf{H}}_{m,\ell,1:r,1:r}); \text{vec}(\hat{\mathbf{H}}_{m,\ell,r+1:k_1,1:r}); \text{vec}(\hat{\mathbf{H}}_{m,\ell,1:r,r+1:k_2}); \\ \boldsymbol{\theta}_{\ell,k+1:p} &= \text{vec}(\hat{\mathbf{H}}_{m,\ell,r+1:k_1,r+1:k_2}).\end{aligned}\quad (10)$$

This is shown in step 14 of Algorithm 2. Now we proceed similarly to Section 2.2. We construct a per-task optimal design for the rotated arm set $\underline{\mathcal{W}}_m, \underline{\mathcal{G}}_m$ and define the $\underline{\mathbf{w}} = [\tilde{\mathbf{g}}_{m,1:d_1}; \tilde{\mathbf{v}}_{m,1:d_2}]$ and $\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}_m$ where $\tilde{\mathbf{g}}_m \in \mathcal{G}_m$ and $\tilde{\mathbf{v}}_m \in \mathcal{V}_m$ respectively. Following (5) we know that to minimize the sample complexity for the m -th bilinear bandit we need to sample according to G -optimal design

$$\hat{\mathbf{b}}_{m,\ell}^G = \arg \min_{\mathbf{b}_{m,\underline{\mathbf{w}}}} \max_{\mathbf{w}, \mathbf{w}' \in \underline{\mathcal{W}}_{m,\ell}} \|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|^2_{(\sum_{\mathbf{w} \in \underline{\mathcal{W}}_m} \mathbf{b}_{m,\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda_{m,\ell}/n)^{-1}} \quad (11)$$

Then **GOBLIN** runs G -optimal design on the arm set $\underline{\mathcal{W}}_\ell$ following the (11) and then samples each $\underline{\mathbf{w}} \in \underline{\mathcal{W}}_{m,\ell}$ for $\lceil \tau_{m,\ell}^G \hat{\mathbf{b}}_{m,\ell}^G \rceil$ times where $\hat{\mathbf{b}}_{m,\ell}^G$ is the solution to the G -optimal design, and τ_ℓ^G is defined in step 17 of Algorithm 2. So the total length of phase ℓ , combining stages 1, 2 and 3 is $(\tau_\ell^E + \sum_m \tilde{\tau}_{m,\ell}^E + \sum_m \tau_{m,\ell}^G)$ rounds. Observe that the stage 1 and 2 design is on the whole arm set $\tilde{\mathcal{W}}, \mathcal{W}_m$ whereas the stage 3 design is on the refined active set $\underline{\mathcal{W}}_{m,\ell}$. Let at the stage 3 of ℓ -th phase the actions sampled be denoted by the matrix $\underline{\mathbf{W}}_{m,\ell} \in \mathbb{R}^{\tau_{m,\ell}^G \times k_1 k_2}$ and observed rewards $\mathbf{r}_m \in \mathbb{R}^{\tau_{m,\ell}^G \times k_1 k_2}$. Define the positive diagonal matrix $\Lambda_{m,\ell}$ according to (6) but set $p = k_1 k_2$ and $k = (k_1 + k_2)r$. Then similar to Section 2.2 we can build for each task m only from the observations from this phase

$$\hat{\boldsymbol{\theta}}_{m,\ell} = \arg \min_{\boldsymbol{\theta}} \frac{1}{2} \|\underline{\mathbf{W}}_{m,\ell} \boldsymbol{\theta} - \mathbf{r}_m\|_2^2 + \frac{1}{2} \|\boldsymbol{\theta}\|_{\Lambda_{m,\ell}}^2 \quad (12)$$

Finally **GOBLIN** eliminates the sub-optimal arms using the estimator $\hat{\boldsymbol{\theta}}_{m,\ell}$ to build the next phase active set $\underline{\mathcal{W}}_{m,\ell}$ and stops when $|\underline{\mathcal{W}}_{m,\ell}| = 1$. The full pseudo-code is given in Algorithm 2.

3.4 Sample Complexity analysis of Multi-task **GOBLIN**

We now present the sample complexity of **GOBLIN** for the multi-task setting.

Theorem 2. (informal) *With probability at least $1 - \delta$, **GOBLIN** returns the best arms $\mathbf{x}_{m,*}, \mathbf{z}_{m,*}$ for each task m , and the total number of samples is bounded by $\tilde{O}\left(\frac{M(k_1+k_2)r}{\Delta^2} + \frac{M\sqrt{k_1 k_2} r}{S_r} + \frac{\sqrt{d_1 d_2} r}{S_r}\right)$.*

Discussion 2. In Theorem 2 the first quantity is the sample complexity to identify the best arms $\mathbf{x}_{m,*}, \mathbf{z}_{m,*}$ and the second quantity is the number of samples to learn $\mathbf{S}_{m,*}$ for each task m . This is required to rotate the arms to reach the effective dimension of $(k_1 + k_2)r$. Finally, the third quantity is the number of samples needed to learn $\Theta_{m,*}$ (which in turn is used to estimate the feature extractors \mathbf{B}_1 and \mathbf{B}_2 to learn the $\mathbf{S}_{m,*}$). Again we assume that $S_r = \Theta(1/\sqrt{r})$ (Kang et al., 2022). So the sample complexity of multi-task **GOBLIN** scales as $\tilde{O}(M(k_1 + k_2)r/\Delta^2)$. However, if one runs DouExpDes (Du et al., 2023) then the sample complexity will scale as $\tilde{O}(M(k_1 k_2)/\Delta^2)$ which is worse than **GOBLIN** when $r \ll k_1$ or k_2 .

Algorithm 2 G-Optimal Design for Bilinear Bandits (**GOBLIN**) for multi-task setting

- 1: Input: arm set \mathcal{X}, \mathcal{Z} , confidence δ , rank r of Θ_* , spectral bound S_r of Θ_* , $S, S_{m,\ell}^\perp = \frac{8k_1 k_2 r}{\tau_{m,\ell}^E S_r^2} \log(\frac{k_1+k_2}{\delta_\ell})$, $\lambda, \lambda_{m,\ell}^\perp = \frac{\tau_{m,\ell}^{G-1}}{(8(k_1+k_2)r \log(1+\tau_{m,\ell-1}^G/\lambda))}$. Let $p = k_1 k_2$, $k = (k_1 + k_2)r$.
 - 2: Let $\mathcal{W}_{m,1} \leftarrow \mathcal{W}_m, \ell \leftarrow 1$, $\tau_0^G = \log(4\ell^2 |\mathcal{X}|/\delta)$. Define $\Lambda_{m,\ell}$ as in (6), $B_{m,*}^\ell := (8\sqrt{\lambda}S + \sqrt{\lambda_{m,\ell}^\perp S_{m,\ell}^\perp})$.
 - 3: Define arm $\bar{\mathbf{w}} = [\mathbf{x}_{1:d_1}; \mathbf{z}_{1:d_2}]$ and $\bar{\mathbf{w}} \in \bar{\mathcal{W}}_m$. Let $\tau_\ell^E = \frac{\sqrt{8d_1 d_2 r \log(4\ell^2 |\mathcal{W}|/\delta_\ell)}}{S_r}$. Let E -optimal design be $\mathbf{b}_\ell^E = \arg \min_{\mathbf{b} \in \Delta_{\bar{\mathcal{W}}}} \|(\sum_{\bar{\mathbf{w}} \in \bar{\mathcal{W}}} \mathbf{b}_{\bar{\mathbf{w}}} \bar{\mathbf{w}} \bar{\mathbf{w}}^\top)^{-1}\|$.
 - 4: **while** $\exists m \in [M], |\mathcal{W}_{m,\ell}| > 1$ **do**
 - 5: $\epsilon_\ell = 2^{-\ell}$, $\delta_\ell = \delta/\ell^2$.
 - 6: **(Stage 1:) Explore the Low-Rank Subspace**
 - 7: Pull arm $\bar{\mathbf{w}} \in \bar{\mathcal{W}}$ exactly $\lceil \hat{\mathbf{b}}_{\bar{\mathbf{w}}}^E \tau_\ell^E \rceil$ times for each task m and observe rewards $\{r_{m,t}\}_{t=1}^{\tau_\ell^E}$.
 - 8: Compute $\hat{\mathbf{Z}}_\ell$ using (8).
 - 9: **(Stage 2:) Build $\hat{\mathbf{S}}_{m,\ell}$ for each task m**
 - 10: Let $\hat{\mathbf{B}}_{1,\ell}, \hat{\mathbf{B}}_{2,\ell}$ be the top- k_1 left and top- k_2 right singular vectors of $\hat{\mathbf{Z}}_\ell$ respectively. Build $\tilde{\mathbf{g}}_m = \mathbf{x}^\top \hat{\mathbf{B}}_{1,\ell}$ and $\tilde{\mathbf{v}}_m = \mathbf{z}^\top \hat{\mathbf{B}}_{2,\ell}$ for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$ for each m .
 - 11: Define a vectorized arm $\tilde{\mathbf{w}} = [\tilde{\mathbf{g}}_{m,1:k_1}; \tilde{\mathbf{v}}_{m,1:k_2}]$ and $\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}_m$ for each m . Let $\tilde{\tau}_{m,\ell}^E = \frac{\sqrt{8k_1 k_2 r \log(4\ell^2 |\mathcal{W}|/\delta_\ell)}}{S_r}$, and $\tilde{\mathbf{b}}_{m,\ell}^E = \arg \min_{\mathbf{b}_m \in \Delta_{\tilde{\mathcal{W}}_m}} \|(\sum_{\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}_m} \mathbf{b}_{m,\tilde{\mathbf{w}}} \tilde{\mathbf{w}} \tilde{\mathbf{w}}^\top)^{-1}\|$.
 - 12: Pull arm $\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}_m$ exactly $\lceil \tilde{\mathbf{b}}_{m,\ell}^E \tilde{\tau}_{m,\ell}^E \rceil$ times and observe rewards $r_{m,t}$, for $t = 1, \dots, \tilde{\tau}_{m,\ell}^E$, for each task m . Then compute $\hat{\mathbf{S}}_{m,\ell}$ using (9) for each m .
 - 13: **(Stage 3:) Reduction to low dimensional linear bandits for each task m**
 - 14: SVD of $\hat{\mathbf{S}}_{m,\ell} = \hat{\mathbf{U}}_{m,\ell} \hat{\mathbf{D}}_{m,\ell} \hat{\mathbf{V}}_{m,\ell}^\top$. Rotate arms in active set $\mathcal{W}_{m,\ell-1}$ to build $\mathcal{W}_{m,\ell}$ using (10).
 - 15: Let $\hat{\mathbf{b}}_{m,\ell}^G = \arg \min_{\mathbf{b}_m, \mathbf{w}} \max_{\mathbf{w}, \mathbf{w}' \in \mathcal{W}_{m,\ell}} \|\mathbf{w} - \mathbf{w}'\|^2_{(\sum_{\mathbf{w}_m \in \mathcal{W}_m} \mathbf{b}_{m,\mathbf{w}_m} \mathbf{w}_m \mathbf{w}_m^\top + \Lambda_{m,\ell}/n)^{-1}}$.
 - 16: Define $\rho^G(\mathcal{Y}(\mathcal{W}_{m,\ell})) = \min_{\mathbf{b}_m, \mathbf{w}} \max_{\mathbf{w}, \mathbf{w}' \in \mathcal{W}_{m,\ell}} \|\mathbf{w} - \mathbf{w}'\|^2_{(\sum_{\mathbf{w} \in \mathcal{W}_m} \mathbf{b}_{m,\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \frac{\Lambda_{m,\ell}}{n})^{-1}}$.
 - 17: Set $\tau_{m,\ell}^G = \frac{64B_{m,*}^\ell \rho^G(\mathcal{Y}(\mathcal{W}_{m,\ell})) \log(4\ell^2 |\mathcal{W}_m|/\delta_\ell)}{\epsilon_\ell^2}$. Then pull arm $\mathbf{w} \in \mathcal{W}_m$ for each task m exactly $\lceil \hat{\mathbf{b}}_{m,\ell}^G \tau_{m,\ell}^G \rceil$ times and construct the least squares estimator $\hat{\boldsymbol{\theta}}_{m,\ell}$ using only the observations of this phase where $\hat{\boldsymbol{\theta}}_{m,\ell}$ is defined in (12).
 - 18: Eliminate arms such that $\mathcal{W}_{m,\ell+1} \leftarrow \mathcal{W}_{m,\ell} \setminus \left\{ \mathbf{w}_m \in \mathcal{W}_{m,\ell} : \max_{\mathbf{w}'_m \in \mathcal{W}_{m,\ell}} \langle \mathbf{w}'_m - \mathbf{w}_m, \hat{\boldsymbol{\theta}}_{m,\ell} \rangle > 2\epsilon_{m,\ell} \right\}$
 - 19: $\ell \leftarrow \ell + 1$
 - 20: Output the arm in $\mathcal{W}_{m,\ell}$ and reshape to get the $\hat{\mathbf{x}}_{m,*}$ and $\hat{\mathbf{z}}_{m,*}$ for each task m .
-

Proof (Overview) of Theorem 2: Step 1 (Subspace estimation in high dimension): The first steps diverge from the proof technique of Theorem 1. We now build the average estimator $\hat{\mathbf{Z}}_\ell$ to estimate the quantity $\mathbf{Z}_* = \frac{1}{M} \sum_{m=1}^M \Theta_{*,m}$ using (8). This requires us to modify the Lemma 3 in Appendix A.1 and apply Stein's lemma (Lemma 1) to get a bound of $\|\hat{\mathbf{Z}}_\ell - \mathbf{Z}_*\|_F^2 \leq \frac{C_1 d_1 d_2 r \log(2(d_1+d_2)/\delta)}{\tau_\ell^E}$ for some $C_1 > 0$. This is shown in Lemma 12 in Appendix A.6. Summing up over $\ell = 1$ to $\lceil \log_2(4\Delta^{-1}) \rceil$ we get that the total samples complexity of the first stage is bounded by $\tilde{O}(\sqrt{d_1 d_2 r}/S_r)$.

Step 2 (Estimation of left and right feature extractors): Now using the estimator in (8) we get a good estimation of the feature extractors \mathbf{B}_1 and \mathbf{B}_2 . Let $\hat{\mathbf{B}}_{1,\ell}, \hat{\mathbf{B}}_{2,\ell}$ be the top- k_1 left and top- k_2 right singular vectors of $\hat{\mathbf{Z}}_\ell$ respectively. Then using the Davis-Kahan sin θ Theorem (Bhatia, 2013) in Lemma 14, 15 (Appendix A.6) we have $\|(\hat{\mathbf{B}}_{1,\ell}^\perp)^\top \mathbf{B}_1\|, \|(\hat{\mathbf{B}}_{2,\ell}^\perp)^\top \mathbf{B}_2\| \leq \tilde{O}(\sqrt{(d_1 + d_2)r/M\tau_\ell^E})$.

Step 3 (Estimation of $\hat{\mathbf{S}}_{m,\ell}$ in low dimension): Now we estimate the quantity $\hat{\mathbf{S}}_{m,\ell} \in \mathbb{R}^{k_1 \times k_2}$ for each task m . To do this we first build the latent arms $\tilde{\mathbf{g}}_m = \mathbf{x}^\top \hat{\mathbf{U}}_\ell$ and $\tilde{\mathbf{v}}_m = \mathbf{z}^\top \hat{\mathbf{V}}_\ell$ for all $\mathbf{x} \in \mathcal{X}$

and $\mathbf{z} \in \mathcal{Z}$ for each m , and sample them following the E -optimal design in step 12 of Algorithm 2. We also show in Lemma 16 (Appendix A.6) that $\sigma_{\min}(\sum_{\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}} \mathbf{b}_{\tilde{\mathbf{w}}} \tilde{\mathbf{w}} \tilde{\mathbf{w}}^\top) > 0$ which enables us to sample following E -optimal design. Then use the estimator in (9). Then in Lemma 19 we show that $\|\hat{\mathbf{S}}_{m,\ell} - \mu^* \mathbf{S}_{m,*}\|_F^2 \leq C_1 k_1 k_2 r \log\left(\frac{2(k_1+k_2)}{\delta_\ell}\right) / \tau_{m,\ell}^E$ holds with probability greater than $(1 - \delta)$. Also, note that in the second phase by setting $\tilde{\tau}_{m,\ell}^E = \sqrt{8k_1 k_2 r \log(4\ell^2 |\mathcal{W}| / \delta_\ell)} / S_r$ and sampling each arm $\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}$ exactly $\lceil \hat{\mathbf{b}}_{\ell, \tilde{\mathbf{w}}}^E \tilde{\tau}_{m,\ell}^E \rceil$ times we are guaranteed that $\|\theta_{k+1:p}^*\|_2 = O(k_1 k_2 r / \tilde{\tau}_{m,\ell}^E)$ in the ℓ -th phase. Summing up over $\ell = 1$ to $\lceil \log_2(4\Delta^{-1}) \rceil$ across each task M we get that the total samples complexity of the second stage is bounded by $\tilde{O}(M \sqrt{k_1 k_2 r} / S_r)$.

Step 4 (Convert to $k_1 k_2$ bilinear bandits): Once GOBLIN recovers $\hat{\mathbf{S}}_{m,\tau_\ell^E}$ it rotates the arm set following (10) to build \mathcal{W}_m to get the $k_1 k_2$ bilinear bandits. The rest of the steps follow the same way as in steps 2, 3 and 4 of proof of Theorem 1.

Step 5 (Total Samples): We show the total samples in the third phase are bounded by $O(\frac{k}{\gamma_y^2} \log(\frac{k \log_2(\Delta^{-1}) |\mathcal{W}|}{\delta}) \lceil \log_2(\Delta^{-1}) \rceil)$ where the effective dimension $k = (k_1 + k_2)r$. The total samples of phase ℓ is given by $\tau_\ell^E + \sum_m (\tilde{\tau}_{m,\ell}^E + \tau_{m,\ell}^G)$. Finally, we get the total sample complexity by summing over all phases from $\ell = 1$ to $\lceil \log_2(4\Delta^{-1}) \rceil$. The claim of the theorem follows by noting $\tilde{O}(k/\gamma_y^2) \leq \tilde{O}(k/\Delta^2)$.

4 Experiments

In this section, we conduct proof-of-concept experiments on both single and multi-task bilinear bandits. In the single-task experiment, we compare against the state-of-the-art RAGE algorithm (Fiez et al., 2019). We show in Figure 1 (left) that GOBLIN requires fewer samples than the RAGE with an increasing number of arms. In the multi-task experiment, we compare against the state-of-the-art DouExpDes algorithm (Du et al., 2023). We show in Figure 1 (right) that GOBLIN requires fewer samples than DouExpDes with an increasing number of tasks. As experiments are not a central contribution, we defer a fuller description of the experimental set-up to Appendix A.8.




Figure 1: (Left) Single-task experiment: results show the number of samples required to identify the optimal action pair for differing numbers of actions. (Right) Multi-task experiment: results show the number of samples required to identify the optimal action pair for varying numbers of tasks.

5 Conclusions and Future Directions

In this paper, we formulated the first pure exploration multi-task representation learning problem. We introduce an algorithm, GOBLIN that achieves a sample complexity bound of $\tilde{O}((d_1 + d_2)r/\Delta^2)$ which improves upon the $\tilde{O}(d_1 d_2)/\Delta^2$ sample complexity of RAGE (Fiez et al., 2019) in a single-task setting. We then extend GOBLIN for multi-task pure exploration bilinear bandit problems by learning latent features which enables sample complexity that scales as $\tilde{O}(M(k_1 + k_2)r/\Delta^2)$ which improves over the $\tilde{O}(M(k_1 k_2)/\Delta^2)$ sample complexity of DouExpDes (Du et al., 2023). Our analysis opens an exciting opportunity to analyze representation learning in the kernel and neural bandits (Zhu et al., 2021; Mason et al., 2021). We can leverage the fact that this type of optimal design does not require the arm set to be an ellipsoid (Du et al., 2023) which enables us to extend our analysis to non-linear representations.

References

- Agarwal, D., Chen, B.-C., and Elango, P. (2009). Explore/exploit schemes for web content optimization. In *2009 Ninth IEEE International Conference on Data Mining*, pages 1–10. IEEE.
- Bengio, Y., Courville, A., and Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828.
- Bhatia, R. (2013). *Matrix analysis*, volume 169. Springer Science & Business Media.
- Bragman, F. J., Tanno, R., Eaton-Rosen, Z., Li, W., Hawkes, D. J., Ourselin, S., Alexander, D. C., McClelland, J. R., and Cardoso, M. J. (2018). Uncertainty in multitask learning: joint representations for probabilistic mr-only radiotherapy planning. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part IV 11*, pages 3–11. Springer.
- Du, S. S., Hu, W., Kakade, S. M., Lee, J. D., and Lei, Q. (2020). Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*.
- Du, Y., Huang, L., and Sun, W. (2023). Multi-task representation learning for pure exploration in linear bandits. *arXiv preprint arXiv:2302.04441*.
- Fiez, T., Jain, L., Jamieson, K. G., and Ratliff, L. (2019). Sequential experimental design for transductive linear bandits. *Advances in neural information processing systems*, 32.
- Jamieson, K. and Jain, L. (2022). Interactive machine learning.
- Jun, K.-S., Willett, R., Wright, S., and Nowak, R. (2019). Bilinear bandits with low-rank structure. In *International Conference on Machine Learning*, pages 3163–3172. PMLR.
- Kang, Y., Hsieh, C.-J., and Lee, T. C. M. (2022). Efficient frameworks for generalized low-rank matrix bandit problems. *Advances in Neural Information Processing Systems*, 35:19971–19983.
- Katz-Samuels, J., Jain, L., Jamieson, K. G., et al. (2020). An empirical process approach to the union bound: Practical algorithms for combinatorial and linear bandits. *Advances in Neural Information Processing Systems*, 33:10371–10382.
- Kiefer, J. and Wolfowitz, J. (1960). The equivalence of two extremum problems. *Canadian Journal of Mathematics*, 12:363–366.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit algorithms*. Cambridge University Press.
- Li, J., Zhang, H., Zhang, L., Huang, X., and Zhang, L. (2014). Joint collaborative representation with multitask learning for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 52(9):5923–5936.
- Lu, Y., Meisami, A., and Tewari, A. (2021). Low-rank generalized linear bandit problems. In *International Conference on Artificial Intelligence and Statistics*, pages 460–468. PMLR.
- Luo, Y., Zhao, X., Zhou, J., Yang, J., Zhang, Y., Kuang, W., Peng, J., Chen, L., and Zeng, J. (2017). A network integration approach for drug-target interaction prediction and computational drug repositioning from heterogeneous information. *Nature communications*, 8(1):573.
- Mason, B., Camilleri, R., Mukherjee, S., Jamieson, K., Nowak, R., and Jain, L. (2021). Nearly optimal algorithms for level set estimation. *arXiv preprint arXiv:2111.01768*.
- Maurer, A., Pontil, M., and Romera-Paredes, B. (2016). The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32.
- Minsker, S. (2018). Sub-gaussian estimators of the mean of a random matrix with heavy-tailed entries. *The Annals of Statistics*, 46(6A):2871–2903.
- Mukherjee, S., Tripathy, A. S., and Nowak, R. (2022). Chernoff sampling for active testing and extension to active regression. In *International Conference on Artificial Intelligence and Statistics*, pages 7384–7432. PMLR.

- Mukherjee, S., Xie, Q., Hanna, J., and Nowak, R. (2023). Speed: Experimental design for policy evaluation in linear heteroscedastic bandits. *arXiv preprint arXiv:2301.12357*.
- Pukelsheim, F. (2006). *Optimal design of experiments*. SIAM.
- Reyes, L. J. P., Oviedo, N. B., Camacho, E. C., and Calderon, J. M. (2021). Adaptable recommendation system for outfit selection with deep learning approach. *IFAC-PapersOnLine*, 54(13):605–610.
- Rockafellar, R. (2015). Convex analysis. princeton landmarks in mathematics and physics.
- Shamir, O. (2011). A variant of azuma’s inequality for martingales with subgaussian tails. *arXiv preprint arXiv:1110.2392*.
- Shen, Q., Han, S., Han, Y., and Chen, X. (2023). User review analysis of dating apps based on text mining. *Plos one*, 18(4):e0283896.
- Soare, M., Lazaric, A., and Munos, R. (2014). Best-arm identification in linear bandits. *Advances in Neural Information Processing Systems*, 27.
- Stein, C., Diaconis, P., Holmes, S., and Reinert, G. (2004). Use of exchangeable pairs in the analysis of simulations. *Lecture Notes-Monograph Series*, pages 1–26.
- Tripuraneni, N., Jin, C., and Jordan, M. (2021). Provable meta-learning of linear representations. In *International Conference on Machine Learning*, pages 10434–10443. PMLR.
- Valko, M., Munos, R., Kveton, B., and Kocák, T. (2014). Spectral bandits for smooth graph functions. In *International Conference on Machine Learning*, pages 46–54. PMLR.
- Yang, J., Hu, W., Lee, J. D., and Du, S. S. (2020). Impact of representation learning in linear bandits. *arXiv preprint arXiv:2010.06531*.
- Yang, J., Lei, Q., Lee, J. D., and Du, S. S. (2022). Nearly minimax algorithms for linear bandits with shared representation. *arXiv preprint arXiv:2203.15664*.
- Zhang, D., Shen, D., Initiative, A. D. N., et al. (2012). Multi-modal multi-task learning for joint prediction of multiple regression and classification variables in alzheimer’s disease. *NeuroImage*, 59(2):895–907.
- Zhao, Y., Kosorok, M. R., and Zeng, D. (2009). Reinforcement learning design for cancer clinical trials. *Statistics in medicine*, 28(26):3294–3315.
- Zhu, Y., Zhou, D., Jiang, R., Gu, Q., Willett, R., and Nowak, R. (2021). Pure exploration in kernel and neural bandits. *Advances in neural information processing systems*, 34:11618–11630.

A Appendix

A.1 Probability Tools and Previous Results

In this section, we state useful lemmas we use in our proofs and previous results.

Lemma 1. (*Generalized Stein's Lemma, (Stein et al., 2004)*) For a random variable X with continuously differentiable density function $\mathbf{p} : \mathbb{R}^d \rightarrow \mathbb{R}$, and any continuously differentiable function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Let $Q(\cdot)$ be a scoring function defined in Definition 1. If the expected values of both $\nabla f(X)$ and $f(X) \cdot Q(X)$ regarding the density p exist, then they are identical, i.e.

$$\mathbb{E}[f(X) \cdot Q(X)] = \mathbb{E}[\nabla f(X)].$$

Lemma 2. (*Minsker, 2018*) Define $\|\mathbf{A}\|_{op}$ as the operator norm of \mathbf{A} . Let $\mathbf{Y}_1, \dots, \mathbf{Y}_n \in \mathbb{R}^{d_1 \times d_2}$ be a sequence of independent real random matrices, and assume that

$$\sigma_n^2 \geq \max \left(\left\| \sum_{j=1}^n \mathbb{E}(\mathbf{Y}_j \mathbf{Y}_j^\top) \right\|_{op}, \left\| \sum_{j=1}^n \mathbb{E}(\mathbf{Y}_j^\top \mathbf{Y}_j) \right\|_{op} \right).$$

Then for any $t \in \mathbb{R}^+$ and $\nu \in \mathbb{R}^+$, it holds that,

$$P \left(\left\| \sum_{j=1}^n \tilde{\psi}_\nu(\mathbf{Y}_j) - \sum_{j=1}^n \mathbb{E}(\mathbf{Y}_j) \right\|_{op} \geq t\sqrt{n} \right) \leq 2(d_1 + d_2) \exp \left(\nu t\sqrt{n} + \frac{\nu^2 \sigma_n^2}{2} \right)$$

Lemma 3. (*Restatement of Theorem 4.1 in Kang et al. (2022)*) For any low-rank linear model with samples $\mathbf{X}_1, \dots, \mathbf{X}_{n_1}$ drawn from \mathcal{X} according to \mathcal{D} then for the optimal solution to the nuclear norm regularization problem in (2) with $\nu = \sqrt{2 \log(2(d_1 + d_2)/\delta) / ((4 + S_0^2) M n_1 d_1 d_2)}$ and

$$\gamma_{n_1} = 4 \sqrt{\frac{2(4 + S_0^2) C d_1 d_2 \log(2(d_1 + d_2)/\delta)}{n_1}},$$

with probability at least $1 - \delta$ it holds that:

$$\left\| \hat{\boldsymbol{\Theta}} - \mu^* \boldsymbol{\Theta}_* \right\|_F^2 \leq \frac{C_1 d_1 d_2 r \log \left(\frac{2(d_1 + d_2)}{\delta} \right)}{n_1},$$

for $C_1 = 36(4 + S_0^2)C$, $\|\mathbf{X}\|_F, \|\boldsymbol{\Theta}_*\|_F \leq S_0$, some nonzero constant μ^* , and $\mathbb{E} \left[(S^p(\mathbf{X}))_{ij}^2 \right] \leq C, \forall i, j$.

A.2 G-optimal design on rotated arms

Remark 1. (*G-optimal design on rotated arms:*) Using the concentration inequality in Proposition 1 we can show that for any arbitrary vector $\mathbf{y} \in \mathbb{R}^p$:

$$|\mathbf{y}^\top (\hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}_*)| \leq \|\mathbf{y}\|_{\mathbf{V}_\ell^{-1}} 2\sqrt{14 \log(2/\delta)} + \|\boldsymbol{\theta}_*\|_{\boldsymbol{\Lambda}_\ell} \leq \|\mathbf{y}\|_{\mathbf{V}_\ell^{-1}} \sqrt{8B_*^\ell \log(2/\delta)}$$

where the co-variance matrix $\mathbf{V}_\ell := \sum_{s=1}^{\tau_\ell^E} \mathbf{w}_s \mathbf{w}_s^\top + \boldsymbol{\Lambda}_\ell$ and B_*^ℓ is defined in Proposition 1. Now we want this to hold for all $\mathbf{y} \in \mathcal{Y}^*(\mathcal{W}_\ell)$, and so we need to union bound over $\mathcal{W} \supseteq \mathcal{W}_\ell$ replacing δ with $\delta/|\mathcal{W}|$. Set the phase length $\tau_\ell^G := \left\lceil \frac{64B_*^\ell \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \log(4\ell^2 |\mathcal{W}|/\delta)}{\epsilon_\ell^2} \right\rceil$ where $\rho^G(\mathcal{Y}(\mathcal{W}_\ell))$ is defined in step 14 of Algorithm 1.

Then for the allocation $2\lfloor \mathbf{b}_{\mathbf{w}}^G \tau_\ell^G \rfloor$ for each $\mathbf{b}_{\mathbf{w}} \in \mathcal{W}_\ell$, we have for each $\mathbf{w} \in \mathcal{W}_\ell \setminus \mathbf{w}_*$ that with probability at least $1 - \delta$,

$$(\mathbf{w}_* - \mathbf{w})^\top \hat{\boldsymbol{\theta}}_\ell \geq (\mathbf{w}_* - \mathbf{w})^\top \boldsymbol{\theta}_* - \|\mathbf{w}_* - \mathbf{w}\|_{(\sum_{\mathbf{w} \in \mathcal{W}} \lceil 2\tau_\ell^G \mathbf{b}_{\mathbf{w}}^* \rceil \mathbf{w} \mathbf{w}^\top + 2\tau_\ell^G \boldsymbol{\Lambda}_\ell / \tau_\ell^G)^{-1}} \sqrt{8B_*^\ell \log(2|\mathcal{W}|/\delta)}$$

since for every $\underline{\mathbf{w}}_* - \underline{\mathbf{w}} \in \mathcal{Y}^*(\mathcal{W})$ we have

$$\begin{aligned} (\underline{\mathbf{w}}_* - \underline{\mathbf{w}})^\top \left(2 \sum_{\underline{\mathbf{w}} \in \mathcal{W}} [\tau_\ell^G \mathbf{b}_{\underline{\mathbf{w}}}^*] \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + 2 \frac{\tau_\ell^G \Lambda_\ell}{\tau_\ell^G} \right)^{-1} (\underline{\mathbf{w}}_* - \underline{\mathbf{w}}) &\leq \frac{1}{\tau_\ell^G} \|\underline{\mathbf{w}}_* - \underline{\mathbf{w}}\|^2_{(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}}^* \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \frac{\Lambda_\ell}{\tau_\ell^G})^{-1}} \\ &\leq \frac{((\underline{\mathbf{w}}_* - \underline{\mathbf{w}})^\top \boldsymbol{\theta}_*)^2}{\sqrt{8B_*^\ell \log(2|\mathcal{W}|/\delta)}}. \end{aligned}$$

The last inequality follows by plugging in the value of τ_ℓ^G and $\rho^G(\mathcal{Y}(\mathcal{W}_\ell))$. Hence to minimize the number of samples τ_ℓ^G in phase ℓ we can re-arrange the above equation to show that

$$\tau_\ell^G \geq \sqrt{8B_*^\ell \log(2|\mathcal{W}|/\delta)} \max_{\underline{\mathbf{w}} \in \mathcal{W}_\ell \setminus \underline{\mathbf{w}}_*} \frac{\|\underline{\mathbf{w}}_* - \underline{\mathbf{w}}\|^2_{(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}}^* \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda/n)^{-1}}}{(\underline{\mathbf{w}}_* - \underline{\mathbf{w}})^\top \boldsymbol{\theta}_*}$$

Hence, to minimize the sample complexity for the bilinear setting we need to sample according to

$$\mathbf{b}_\ell^G = \arg \min_{\mathbf{b}} \max_{\underline{\mathbf{w}}} \frac{\|\underline{\mathbf{w}}_* - \underline{\mathbf{w}}\|^2_{(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda/n)^{-1}}}{(\underline{\mathbf{w}}_* - \underline{\mathbf{w}})^\top \boldsymbol{\theta}_*} \quad (13)$$

However, note that we do know the identity of $\underline{\mathbf{w}}_*$ or the gaps $(\underline{\mathbf{w}}_* - \underline{\mathbf{w}})^\top \boldsymbol{\theta}_*$. So we replace the gaps with a lower bound of $\epsilon = 2^{-t}$ and compare against every pair of arms $\underline{\mathbf{w}}$ and $\underline{\mathbf{w}}'$ as follows:

$$\mathbf{b}_\ell^G = \arg \min_{\mathbf{b}_{\underline{\mathbf{w}}}} \max_{\underline{\mathbf{w}}, \underline{\mathbf{w}}' \in \mathcal{W}_\ell} \|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|^2_{(\sum_{\underline{\mathbf{w}} \in \mathcal{W}} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda_\ell/n)^{-1}} \quad (14)$$

This is shown in step 12 of Algorithm 1.

A.3 Application of Stein's Lemma

We also present the following two definitions from Kang et al. (2022) to facilitate analysis via Stein's method:

Definition 1. (Score Function) Let $\mathbf{p} : \mathbb{R} \rightarrow \mathbb{R}$ be a univariate probability density function defined on \mathbb{R} . The score function $Q^{\mathbf{p}} : \mathbb{R} \rightarrow \mathbb{R}$ regarding density $\mathbf{p}(\cdot)$ is defined as:

$$Q^{\mathbf{p}}(x) = -\nabla_x \log(\mathbf{p}(x)) = -\nabla_x \mathbf{p}(x)/\mathbf{p}(x), \quad x \in \mathbb{R}.$$

In particular, for a random matrix with its entrywise probability density $\mathbf{p} = (p_{ij}) : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^{d_1 \times d_2}$, we define its score function $Q^{\mathbf{p}} = (Q_{ij}^{\mathbf{p}}) : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^{d_1 \times d_2}$ as $Q_{ij}^{\mathbf{p}}(x) = Q^{p_{ij}}(x)$ by applying the univariate score function to each entry of \mathbf{p} independently.

Assumption 3. The norm of true parameter $\boldsymbol{\Theta}^*$ and feature matrices in \mathcal{X} is bounded: there exists $S \in \mathbb{R}^+$ such that for all arms $\mathbf{X} \in \mathcal{X}$, $\|\mathbf{X}\|_F, \|\boldsymbol{\Theta}^*\|_F \leq S_0$

Assumption 4. (Finite second-moment score) There exists a sampling distribution \mathcal{D} over \mathcal{X} such that for the random matrix \mathbf{X} drawn from \mathcal{D} with its associated density $\mathbf{p} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^{d_1 \times d_2}$, we have $\mathbb{E} \left[(Q^{\mathbf{p}}(\mathbf{X}))_{ij}^2 \right] \leq C, \forall i, j$

Definition 2. Given a rectangular matrix $\mathbf{A} \in \mathbb{R}^{d_1 \times d_2}$, the (Hermitian) dilation $\mathcal{H} : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^{(d_1+d_2) \times (d_1+d_2)}$ is defined as:

$$\mathcal{H}(\mathbf{A}) = \begin{pmatrix} 0 & \mathbf{A} \\ \mathbf{A}^\top & 0 \end{pmatrix}$$

Definition 3. (The function $\tilde{\psi}_\nu$) To explore the valid subspace of the parameter matrix $\boldsymbol{\Theta}_*$, we define a function $\psi : \mathbb{R} \rightarrow \mathbb{R}$ in (15). Let $\mathcal{H}(\cdot)$ be as defined in Definition 2. Then define $\tilde{\psi}_\nu : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^{d_1 \times d_2}$ as $\tilde{\psi}_\nu(\mathbf{A}) = \psi(\nu \mathcal{H}(\mathbf{A}))_{1:d_1, (d_1+1):(d_1+d_2)}/\nu$ for some parameter $\nu \in \mathbb{R}^+$

$$\psi(x) = \begin{cases} \log(1 + x + x^2/2), & x \geq 0 \\ -\log(1 - x + x^2/2), & x < 0 \end{cases} \quad (15)$$

A.4 Single-task Pure Exploration Proofs

Good Event: Define the good event \mathcal{F}_ℓ in phase ℓ that **GOBLIN** has a good estimate of Θ_* as follows:

$$\mathcal{F}_\ell = \left\| \hat{\Theta}_\ell - \mu^* \Theta_* \right\|_F^2 \leq \frac{C_1 d_1 d_2 r \log \left(\frac{2(d_1 + d_2)}{\delta_\ell} \right)}{\tau_\ell^E}, \quad (16)$$

where, $C_1 = 36(4 + S_0^2)C$, $\|\mathbf{X}\|_F, \|\Theta_*\|_F \leq S_0$, some nonzero constant μ^* , $\mathbb{E} \left[(S^p(\mathbf{X}))_{ij}^2 \right] \leq C, \forall i, j$, and $\hat{\Theta}_\ell$ is the estimate from (2). Then define the good event

$$\mathcal{F} := \bigcap_{\ell=1}^{\infty} \mathcal{F}_\ell. \quad (17)$$

Lemma 4. *The event \mathcal{F} holds with probability greater than $(1 - \delta/2)$.*

Proof. From Lemma 3 we know the event \mathcal{F}_ℓ in (16) holds with probability $(1 - \delta_\ell)$. Taking a union bound over all phases $\ell \geq 1$ and recalling $\delta_\ell := \frac{\delta}{2\ell^2}$, we obtain

$$\begin{aligned} \mathbb{P}(\mathcal{F}) &\geq 1 - \sum_{\ell=1}^{\infty} \mathbb{P}(\mathcal{F}_\ell^c) \\ &\geq 1 - \sum_{\ell=1}^{\infty} \frac{\delta_\ell}{2} \\ &= 1 - \sum_{\ell=1}^{\infty} \frac{\delta}{4\ell^2} \\ &\geq 1 - \frac{\delta}{2}. \end{aligned}$$

This concludes our proof. \square

Now we move to the second stage for the rotated arm set $\underline{\mathbf{w}} \in \underline{\mathcal{W}}$ and prove the following concentration event.

Lemma 5. *For any fixed $\underline{\mathbf{w}} \in \mathbb{R}^p$ and any $\delta > 0$, we have that if $\beta(\theta_*, \delta) = 2\sqrt{14 \log(2/\delta)} + \|\theta_*\|_{\Lambda}$, then at time $\tau_{\ell-1} + 1$ (beginning of phase ℓ):*

$$\mathbb{P} \left(\left| \underline{\mathbf{w}}^\top (\hat{\theta}_\ell - \theta_*) \right| \leq \|\underline{\mathbf{w}}\|_{\mathbf{V}_\ell^{-1}} \beta(\theta_*, \delta) \right) \geq 1 - \delta$$

where, $\mathbf{V}_\ell := \sum_{s=\tau_{\ell-1}+1}^{\tau_\ell} \underline{\mathbf{w}}_s \underline{\mathbf{w}}_s^\top + \Lambda_\ell$.

Proof. We follow the proof technique of Lemma 7 of Valko et al. (2014). Defining $\xi_\ell = \sum_{s=\tau_{\ell-1}+1}^{\tau_\ell} \underline{\mathbf{w}}_s \eta_s$, we have:

$$\begin{aligned} \left| \underline{\mathbf{w}}^\top (\hat{\theta}_\ell - \theta_*) \right| &\stackrel{(a)}{=} \left| \underline{\mathbf{w}}^\top (-\mathbf{V}_\ell^{-1} \Lambda \theta_* + \mathbf{V}_\ell^{-1} \xi_\ell) \right| \\ &\stackrel{(b)}{\leq} \left| \underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \Lambda \theta_* \right| + \left| \underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \xi_\ell \right| \end{aligned} \quad (18)$$

where (a) follows from Woodbury matrix identity and rearranging the terms, and (b) follows from the triangle inequality.

The first term in the right-hand side of (18) is bounded as:

$$\begin{aligned} \left| \underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \Lambda \theta_* \right| &\leq \left\| \underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \Lambda^{1/2} \right\| \left\| \Lambda^{1/2} \theta_* \right\| \\ &\stackrel{(a)}{=} \|\theta_*\|_{\Lambda_\ell} \sqrt{\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \Lambda \mathbf{V}_\ell^{-1} \underline{\mathbf{w}}} \\ &\leq \|\theta_*\|_{\Lambda_\ell} \sqrt{\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \underline{\mathbf{w}}} = \|\theta_*\|_{\Lambda_\ell} \|\underline{\mathbf{w}}\|_{\mathbf{V}_\ell^{-1}} \end{aligned}$$

where, (a) follows as $\|\boldsymbol{\theta}_*\|_{\Lambda_\ell} = \sqrt{\boldsymbol{\theta}_*^\top \Lambda_\ell \boldsymbol{\theta}_*} = \|\Lambda_\ell^{1/2} \boldsymbol{\theta}_*\|$ and similarly for $\|\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \Lambda_\ell^{1/2}\|$. Now consider the second term in the r.h.s. of (18). We have:

$$|\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \xi_\ell| = \left| \sum_{s=\tau_{\ell-1}+1}^{\tau_\ell} (\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \underline{\mathbf{w}}_s) \eta_s \right|.$$

Now note that the arms $(\underline{\mathbf{w}}_s)$ selected by the algorithm during phase ℓ only depend on the proportion \mathbf{b}_*^G (the G-optimal design) and do not depend on the rewards received during the phase $\ell - 1$. Thus, given \mathcal{F}_{j-2} , the sequence $(\underline{\mathbf{w}}_s)_{\tau_{\ell-1}+1 \leq s < \tau_\ell}$ is deterministic. Consequently, one may use a variant of Azuma's inequality (Shamir (2011)) with a 1-sub Gaussian assumption:

$$\mathbb{P} \left(|\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \xi_\ell|^2 \leq 28 \times 2 \log(2/\delta) \times \underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \left(\sum_{s=\tau_{\ell-1}+1}^{\tau_\ell} \underline{\mathbf{w}}_s \underline{\mathbf{w}}_s^\top \right) \mathbf{V}_\ell^{-1} \underline{\mathbf{w}} \mid \mathcal{F}_{\ell-2} \right) \geq 1 - \delta,$$

from which we deduce:

$$\mathbb{P} \left(|\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \xi_\ell|^2 \leq 56 \underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \underline{\mathbf{w}} \log(2/\delta) \mid \mathcal{F}_{\ell-2} \right) \geq 1 - \delta,$$

since $\sum_{s=\tau_{\ell-1}+1}^{\tau_\ell} \underline{\mathbf{w}}_s \underline{\mathbf{w}}_s^\top \prec \mathbf{V}_\ell$. Thus:

$$\mathbb{P} \left(|\underline{\mathbf{w}}^\top \mathbf{V}_\ell^{-1} \xi_\ell| \leq 2 \|\underline{\mathbf{w}}\|_{\mathbf{V}_\ell^{-1}} \sqrt{14 \log(2/\delta)} \right) \geq 1 - \delta$$

Combining everything we get that

$$\mathbb{P} \left(\left| \underline{\mathbf{w}}^\top (\hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}_*) \right| \leq 2 \sqrt{14 \log(2/\delta)} + \|\boldsymbol{\theta}_*\|_{\Lambda_\ell} \right) \leq 1 - \delta.$$

□

We need to change Lemma 6 of Valko et al. (2014) in the following way so that the dependence on horizon n is replaced by $\tau_{\ell-1}^G$. Note that $\tau_{\ell-1}^G$ is the phase length in the $\ell - 1$ -th phase and is determined before the start of phase $\tau_{\ell-1}^G$. Also, note that using the standard analysis of phase-based algorithms in Fiez et al. (2019); Lattimore and Szepesvári (2020) we do not re-use data between phases. First, we need the following support lemma from Valko et al. (2014).

Lemma 6. (Restatement of Lemma 4 from Valko et al. (2014)) Let $\Lambda_\ell = \text{diag}(\lambda_1, \dots, \lambda_\ell^\perp)$ be any diagonal matrix with strictly positive entries. Define $\mathbf{V}_\ell = \Lambda_\ell + \sum_{s=\tau_{\ell-1}^G+1}^{\tau_\ell^G} \underline{\mathbf{w}}_s \underline{\mathbf{w}}_s^\top$. Then for any vectors $(\underline{\mathbf{w}}_s)_{\tau_{\ell-1}^G+1 \leq s \leq \tau_\ell^G}$, such that $\|\underline{\mathbf{w}}_s\|_2 \leq 1$ for all rounds s such that $\tau_{\ell-1}^G + 1 \leq s \leq \tau_\ell^G$, we have that the determinant $|\mathbf{V}_\ell|$ is maximized when all $\underline{\mathbf{w}}_s$ are aligned with the axes.

Lemma 7. Let k be the effective dimension. Then

$$\log \frac{|\mathbf{V}_\ell|}{|\Lambda_\ell|} \leq 8k \log \left(1 + \frac{\tau_{\ell-1}^G}{\lambda} \right)$$

when $\lambda_\ell^\perp = \frac{\tau_{\ell-1}^G}{k \log(1 + \tau_{\ell-1}^G / \lambda)}$.

Proof. We want to bound the determinant $|\mathbf{V}_\ell|$ under the coordinate constraints $\|\underline{\mathbf{w}}_t\|_2 \leq 1$. Let:

$$M(\underline{\mathbf{w}}_1, \dots, \underline{\mathbf{w}}_t) = \left| \Lambda_\ell + \sum_{s=\tau_{\ell-1}^G+1}^{\tau_\ell^G} \underline{\mathbf{w}}_s \underline{\mathbf{w}}_s^\top \right|$$

From Lemma 6 we deduce that the maximum of M is reached when all $\underline{\mathbf{w}}_s$ are aligned with the axes. Let the number of samples of these axes-aligned $\underline{\mathbf{w}}_s$'s during the ℓ -th phase be denoted as

$t_1^\ell, t_2^\ell, \dots, t_p^\ell$ such that $\sum_{i=1}^p t_i^\ell = \tau_\ell^G$. Then we can show that

$$\begin{aligned}
M &\stackrel{(a)}{=} \max_{\mathbf{w}_1, \dots, \mathbf{w}_t, \mathbf{w}_s \in \{\mathbf{e}_1, \dots, \mathbf{e}_p\}} \left| \mathbf{\Lambda}_\ell + \sum_{s=\tau_{\ell-1}^G+1}^{\tau_\ell^G} \mathbf{w}_s \mathbf{w}_s^\top \right| \\
&\stackrel{(b)}{=} \max_{t_1^\ell, \dots, t_p^\ell, \text{ positive integers, } \sum_{i=1}^p t_i^\ell = \tau_\ell^G} |\text{diag}(\lambda_i + t_i^\ell)| \\
&\stackrel{(c)}{\leq} \max_{t_1^\ell, \dots, t_p^\ell, \text{ positive integers, } \sum_{i=1}^p t_i^\ell = \tau_\ell^G} \prod_{i=1}^p (\lambda_i + t_i^\ell)
\end{aligned}$$

where, (a) follows from Lemma 6, (b) follows as the $\text{diag}(\lambda_i + t_i^\ell)$ contains the number of times axis-aligned $\mathbf{w}_s \in \{\mathbf{e}_1, \dots, \mathbf{e}_p\}$ are observed, and (c) follows as the determinant of a diagonal matrix is the product of the diagonal elements. Now we can show that

$$\begin{aligned}
\log \frac{|\mathbf{V}_\ell|}{|\mathbf{\Lambda}_\ell|} &\leq \sum_{i=1}^k \log \left(1 + \frac{t_i^\ell}{\lambda} \right) + \sum_{i=k+1}^p \log \left(1 + \frac{t_i^\ell}{\lambda_i} \right) \\
&\stackrel{(a)}{\leq} k \log \left(1 + \frac{t_i^\ell}{\lambda} \right) + \sum_{i=1}^p \frac{t_i^{\ell-1}}{\lambda_\ell^\perp} \\
&\stackrel{(b)}{\leq} k \log \left(1 + \frac{t_i^\ell}{\lambda} \right) + \frac{\tau_{\ell-1}^G}{\lambda_\ell^\perp} \\
&\stackrel{(c)}{\leq} k \log \left(1 + \frac{t_i^\ell}{\lambda} \right) + k \log \left(1 + \frac{\tau_{\ell-1}^G}{\lambda} \right) \\
&\stackrel{(d)}{\leq} 8k \log \left(1 + \frac{\tau_{\ell-1}^G}{\lambda} \right)
\end{aligned}$$

where, (a) follows as $\log(1 + t_i^\ell/\lambda_i) \leq t_i^{\ell-1}/\lambda_\ell^\perp$, (b) follows as $\sum_{i=1}^p t_i^{\ell-1} = \tau_{\ell-1}^G$, and (c) follows for $\lambda_\ell^\perp = \frac{\tau_{\ell-1}^G}{k \log(1 + \tau_{\ell-1}^G/\lambda)}$ and (d) follows from Lemma 8. \square

Lemma 8. Let $\rho^G(\mathcal{Y}(\mathcal{W}_\ell)) = \min_{\mathbf{w}} \max_{\mathbf{w}', \mathbf{w}'' \in \mathcal{W}_\ell} \|\mathbf{w} - \mathbf{w}'\|_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \mathbf{\Lambda}_\ell/n)^{-1}}$. Recall that $\tau_\ell^G = \frac{8B_*^\ell \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \log(4\ell^2 |\mathcal{W}|/\delta)}{\epsilon_\ell^2}$. Assume $\log(p) \leq k$. Then we can show that

$$\log \left(1 + \frac{\tau_\ell^G}{\lambda} \right) \leq 8k \log \left(1 + \frac{\tau_{\ell-1}^G}{\lambda} \right).$$

Proof. We start by first recalling the definition of

$$\tau_\ell^G = 2\epsilon_\ell^{-2} 8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \log(4\ell^2 |\mathcal{W}|/\delta_\ell).$$

Then we can show the following

$$\begin{aligned}
\frac{\tau_\ell^G}{\tau_{\ell-1}^G} &= \frac{2\epsilon_\ell^{-2} B_*^\ell \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \log(4\ell^2 |\mathcal{W}|/\delta_\ell)}{2\epsilon_{\ell-1}^{-2} B_*^{\ell-1} \rho^G(\mathcal{Y}(\mathcal{W}_{\ell-1})) \log(4(\ell-1)^2 |\mathcal{W}|/\delta_{\ell-1})} \leq \frac{4\rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \left(64(\lambda S^2 + \lambda_\ell^\perp S_\perp^{(2),\ell}) \right)}{\rho^G(\mathcal{Y}(\mathcal{W}_{\ell-1})) \left(64(\lambda S^2 + \lambda_{\ell-1}^\perp S_\perp^{(2),\ell-1}) \right)} \\
&\leq \frac{4\rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \tau_{\ell-1}^G / \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})}{\rho^G(\mathcal{Y}(\mathcal{W}_{\ell-1})) \tau_{\ell-2}^G / \log(1 + \frac{\tau_{\ell-2}^G}{\lambda})} \\
&\stackrel{(a)}{\leq} \frac{4\rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \tau_{\ell-1}^G \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})}{\rho^G(\mathcal{Y}(\mathcal{W}_{\ell-1}))} \\
&\stackrel{(b)}{\leq} \frac{4p}{\gamma_{\mathcal{Y}}^2} \frac{\max_{\mathbf{w} \in \mathcal{W}} \|\mathbf{w}\|_2}{\max_{\mathbf{y} \in \mathcal{Y}(\mathcal{W}_\ell)} \|\mathbf{y}\|_2^2} \tau_{\ell-1}^G \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) = \frac{4p}{C\gamma_{\mathcal{Y}}^2}
\end{aligned}$$

where, (a) follows as $\log(1 + \frac{\tau_{\ell-2}^G}{\lambda}) \geq 1$ and $\log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \geq \log(1 + \frac{\tau_{\ell-2}^G}{\lambda})$. The (b) follows using Lemma 1 from [Fiez et al. \(2019\)](#) such that

$$\max_{\mathbf{y} \in \mathcal{Y}(\mathcal{W}_\ell)} \|\mathbf{y}\|_2^2 / \left(\max_{\mathbf{w} \in \mathcal{W}} \|\mathbf{w}\|_2 \right) \leq \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \leq p/\gamma_{\mathcal{Y}}^2 \xrightarrow{(a_1)} 1 \leq \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \leq p/\gamma_{\mathcal{Y}}^2.$$

where, (a₁) follows as $\|\mathbf{x}\| \leq 1, \|\mathbf{z}\| \leq 1$. This implies that for a constant $C > 0$

$$\begin{aligned} \tau_\ell^G &\leq \frac{4p}{C\gamma_{\mathcal{Y}}^2} (\tau_{\ell-1}^G)^2 \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \implies \log\left(1 + \frac{\tau_\ell^G}{\lambda}\right) \leq \log\left(1 + \frac{4p}{C\gamma_{\mathcal{Y}}^2} (\tau_{\ell-1}^G)^2 \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})\right) \\ &\xrightarrow{(a)} \log\left(1 + \frac{\tau_\ell^G}{\lambda}\right) \leq 4k \log\left(1 + \frac{\tau_{\ell-1}^G}{C\gamma_{\mathcal{Y}}^2 \lambda}\right) \\ &\xrightarrow{(b)} \log\left(1 + \frac{\tau_\ell^G}{\lambda}\right) \leq 8k \log\left(1 + \frac{\tau_{\ell-1}^G}{\lambda}\right) \end{aligned}$$

where, in (a) follows for $\log(p) \leq k$, $\log(a^2 \log(a)) \leq 4 \log(a)$. The (b) follows as $4k \log\left(1 + \frac{\tau_{\ell-1}^G}{C\gamma_{\mathcal{Y}}^2 \lambda}\right) \leq 8k \log\left(1 + \frac{\tau_{\ell-1}^G}{\lambda}\right)$. \square

Lemma 9. The G -optimal design in (5) is equivalent to solving the D -optimal design

$$\mathbf{b}_\ell^D = \arg \max_{\mathbf{b}} \log \frac{\left| \sum_{\mathbf{w} \in \mathcal{W}_\ell} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \mathbf{\Lambda}_\ell \right|}{|\mathbf{\Lambda}_\ell|}.$$

Furthermore, the support of $|\mathbf{b}_\ell^D| \leq \frac{8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})(8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) + 1)}{2}$, where $k = (d_1 + d_2)r$.

Proof. To prove the equivalence between \mathbf{b}_*^G and \mathbf{b}_*^D we need to first show that the regularization matrix $\mathbf{\Lambda}_\ell$ does not depend on \mathbf{w} or $\mathbf{y} = \mathbf{w} - \mathbf{w}'$, where $\mathbf{w} \in \mathcal{W}_\ell$. Define $\text{conv}(\mathcal{X} \cup -\mathcal{X})$ as the convex hull of $\mathcal{X} \cup -\mathcal{X}$. Now recall we have from Lemma 1 of [Fiez et al. \(2019\)](#) that

$$1 \leq \rho^G(\mathcal{Y}(\mathcal{W}_\ell)) \leq p/\gamma_{\mathcal{Y}}^2 \quad (19)$$

where, $\gamma_{\mathcal{Y}} = \max\{c > 0 : c\mathcal{Y} \subset \text{conv}(\mathcal{W}_\ell \cup -\mathcal{W}_\ell)\}$ as the gauge norm of \mathcal{Y} ([Rockafellar, 2015](#)). We can consider the gauge norm $\gamma_{\mathcal{Y}}$ as a problem-dependent constant. Now recall that

$$\begin{aligned} \lambda_\perp^\ell &= \frac{\tau_{\ell-1}^G}{8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})} \leq 2\tau_{\ell-1}^G \leq \frac{64B_*^{\ell-1} \rho^G(\mathcal{Y}(\mathcal{W}_{\ell-1})) \log(4(\ell-1)^2 |\mathcal{W}|/\delta_\ell)}{\epsilon_{\ell-1}^2} \\ &\stackrel{(a)}{\leq} \frac{64(B_*^{\ell-1})^2 p \log(4(\ell-1)^2 |\mathcal{W}|/\delta_\ell)}{\gamma_{\mathcal{Y}}^2 \epsilon_{\ell-1}^2} \\ &\stackrel{(b)}{\leq} \frac{(256(\lambda S^2 + 8p^2 r)) \log(4(\ell-1)^2 p |\mathcal{W}|/\delta_\ell)}{S_r \gamma_{\mathcal{Y}}^2 \epsilon_{\ell-1}^2} \end{aligned}$$

where, (a) follows from (19) and noting that $B_*^{\ell-1} \leq (B_*^{\ell-1})^2$. The (b) follows as $S_\ell^\perp := \frac{8pr}{\tau_\ell^E S_r^2} \log\left(\frac{d_1+d_2}{\delta_\ell}\right)$, $p = d_1 d_2$ and substituting this value and λ_\perp^ℓ in B_*^ℓ we get

$$\begin{aligned} B_*^\ell &\leq (B_*^\ell)^2 \leq \left(256(\lambda S^2 + \lambda_\perp^\ell S_\perp^{(2),\ell})\right) \leq \left(256 \left(\lambda S^2 + \frac{8pr}{S_r^2} \frac{\tau_{\ell-1}^G}{\tau_\ell^E} \cdot \log(4(\ell-1)^2 |\mathcal{W}|/\delta_\ell) \log(d_1 + d_2/\delta_\ell)\right)\right) \\ &\leq \left(256 \left(\lambda S^2 + \frac{8pr}{S_r^2} \rho_{\ell-1}^G(\mathcal{W}) \log(p|\mathcal{W}|/\delta_\ell)\right)\right) \stackrel{(a)}{\leq} (512 (\lambda S^2 + 8p^2 r) / S_r). \end{aligned}$$

Here S_r is the r -th largest eigenvalue of matrix $\mathbf{\Theta}_*$. Substituting this value of λ_ℓ^\perp we can show that $\mathbf{\Lambda}_\ell$ does not depend on \mathbf{w} or $\mathbf{y} = \mathbf{w} - \mathbf{w}'$. The rest of the proof to show equivalence follows the same way as in Theorem 21.1 in [Lattimore and Szepesvári \(2020\)](#).

To bound the support of \mathbf{b}_*^D we proceed as follows: Define the set $\mathcal{Y}(\mathcal{W}_\ell)$ as the set of all arms containing $\mathbf{y} = \underline{\mathbf{w}} - \underline{\mathbf{w}}' \in \mathbb{R}^p$. Then we can use Lemma 7 of [Soare et al. \(2014\)](#) to show that the solution to

$$\max_{\mathbf{y} \in \mathcal{Y}(\mathcal{W}_\ell)} \|\mathbf{y}\|^2 \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1} = \max_{\underline{\mathbf{w}}, \underline{\mathbf{w}}'} \|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|^2 \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1}$$

has a support of atmost $(k_1 + 1)k_1/2$ where $k_1 = 8k \log(1 + \tau_{\ell-1}^G/\lambda)$. The proof follows from the fact that for any pair $(\underline{\mathbf{w}}, \underline{\mathbf{w}}')$ we can show that

$$\|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\| \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1} \leq 2 \max_{\underline{\mathbf{w}}'' \in \mathcal{W}_\ell} \|\underline{\mathbf{w}}''\| \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1}.$$

Then following the work of [Jamieson and Jain \(2022\)](#) Frank-Wolfe algorithm (in section 2.3.1) with the

$$g(\mathbf{b}) = \log \frac{\left| \sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right|}{|\mathbf{\Lambda}_\ell|} = \log \left| \sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right| - \log |\mathbf{\Lambda}_\ell|,$$

and setting for the j -th iteration of the Frank-Wolfe the

$$I_j = \arg \max_{\mathbf{y} \in \mathcal{Y}(\mathcal{W}_\ell)} \|\mathbf{y}\|^2 \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}}^j \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1},$$

and stopping condition

$$\max_{\mathbf{y} \in \mathcal{Y}(\mathcal{W}_\ell)} \|\mathbf{y}\|^2 \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}}^j \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1} \leq 8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \quad (20)$$

This can be done because note that for any $\mathbf{b} \in \Delta_{\mathcal{W}_\ell}$ we have by Kiefer-Wolfowitz Theorem ([Kiefer and Wolfowitz, 1960](#)) that $[\nabla g(\mathbf{b})]_{\mathbf{y}} = \|\mathbf{y}\|^2 \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1} \geq 8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})$. This is because $\mathbf{\Lambda}_\ell$ does not depend on $\underline{\mathbf{w}}$ or \mathbf{y} by the same logic as discussed before. The rest of the proof follows by the same way as in section 2.3.1 in [Jamieson and Jain \(2022\)](#). This will result in a support size of \mathbf{b} at most $8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})(8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) + 1)/2$ following Lemma 7 of [Soare et al. \(2014\)](#). Hence, it follows that solving the Equation (5) will result in a support of $|\mathbf{b}_\ell^D| \leq \frac{8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})(8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) + 1)}{2}$. \square

Proposition 1. *If \mathbf{b}_ℓ^G is the G -optimal design for \mathcal{W}_ℓ then if we pull arm $\underline{\mathbf{w}} \in \mathcal{W}_\ell$ exactly $\lceil \tau^G \mathbf{b}_\ell^G \rceil$ times for some $\tau_\ell^G > 0$ and compute the least squares estimator $\hat{\boldsymbol{\theta}}_\ell$. Then for each $\underline{\mathbf{w}} \in \mathcal{W}_\ell$ we have with probability at least $1 - \delta$*

$$\mathbb{P} \left(\bigcup_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \left\{ \left| \langle \underline{\mathbf{w}}, \hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^* \rangle \right| \leq \sqrt{\frac{64 B_*^\ell k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \log(2|\mathcal{W}|/\delta_\ell)}{\tau_\ell^G}} \right\} \right) \geq 1 - \delta_\ell.$$

where, $\|\boldsymbol{\theta}^*\|_\Lambda \leq \sqrt{\lambda \|\boldsymbol{\theta}_{1:k}\|_2^2 + \lambda_\ell^\perp \|\boldsymbol{\theta}_{k+1:p}\|_2^2} \leq \sqrt{\lambda} S + \sqrt{\lambda_\ell^\perp} S_\perp^\ell$, $B_*^\ell = \sqrt{\lambda} S + \sqrt{\lambda_\ell^\perp} S_\perp^\ell$.

Proof. From Woodbury Matrix Identity we know that for any arbitrary matrix \mathbf{A} and \mathbf{B} , we have the following identity $(\mathbf{A} + \mathbf{B})^{-1} = \mathbf{A}^{-1} - (\mathbf{A} + \mathbf{A}\mathbf{B}^{-1}\mathbf{A})^{-1}$. It follows then that

$$\underline{\mathbf{w}}^\top (\mathbf{A} + \mathbf{B})^{-1} \underline{\mathbf{w}} = \underline{\mathbf{w}}^\top \left(\mathbf{A}^{-1} - (\mathbf{A} + \mathbf{A}\mathbf{B}^{-1}\mathbf{A})^{-1} \right) \underline{\mathbf{w}} \leq \underline{\mathbf{w}}^\top \mathbf{A}^{-1} \underline{\mathbf{w}} = \|\underline{\mathbf{w}}\|_{\mathbf{A}^{-1}}.$$

Hence we can show that,

$$\|\underline{\mathbf{w}}\| \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \lceil \tau_\ell^G \mathbf{b}_\ell^G \rceil \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \mathbf{\Lambda}_\ell \right)^{-1} \leq \underline{\mathbf{w}}^\top \left(\sum_{\underline{\mathbf{w}} \in \mathcal{W}_\ell} \lceil \tau_\ell^G \mathbf{b}_\ell^G \rceil \underline{\mathbf{w}} \underline{\mathbf{w}}^\top \right)^{-1} \underline{\mathbf{w}}. \quad (21)$$

From Lemma 9 we know the support of \mathbf{b}_ℓ^G is less than

$$\frac{8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})(8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) + 1)}{2} \leq (8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}))^2.$$

Also note that $\|\theta^*\|_{\Lambda_\ell} \leq \sqrt{\lambda \|\theta_{1:k}\|_2^2 + \lambda_\perp^\ell \|\theta_{k+1:p}\|_2^2} \leq \sqrt{\lambda}S + \sqrt{\lambda_\perp^\ell S_\perp^\ell}$. Then we can show that

$$\begin{aligned}
\langle \underline{\mathbf{w}}, \hat{\theta}_\ell - \theta^* \rangle &\leq \|\underline{\mathbf{w}}\| \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \lceil \tau_\ell^G \mathbf{b}_\ell^G \rceil \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda \right)^{-1} \left(2\sqrt{14 \log(2/\delta_\ell)} + \sqrt{\lambda}S + \sqrt{\lambda_\perp^\ell S_\perp^\ell} \right) \\
&\leq \frac{1}{\sqrt{\tau_\ell^G}} \|\underline{\mathbf{w}}\| \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \mathbf{b}_{\ell*} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda \right)^{-1} \left(2\sqrt{14 \log(2/\delta_\ell)} + \sqrt{\lambda}S + \sqrt{\lambda_\perp^\ell S_\perp^\ell} \right) \\
&\stackrel{(a)}{\leq} \sqrt{\frac{56 \times 8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \log(2/\delta_\ell)}{\tau_\ell^G}} + \sqrt{\frac{28k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \log(2/\delta_\ell)}{\tau_\ell^G}} (\sqrt{\lambda}S + \sqrt{\lambda_\perp^\ell S_\perp^\ell}) \\
&= \sqrt{\frac{8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \log(2/\delta_\ell)}{\tau_\ell^G}} \left(\sqrt{56} + \underbrace{\sqrt{\lambda}S + \sqrt{\lambda_\perp^\ell S_\perp^\ell}}_{B_*^\ell} \right) \\
&\leq \sqrt{\frac{64B_*^\ell k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \log(2/\delta_\ell)}{\tau_\ell^G}}
\end{aligned}$$

where, (a) follows as $\|\underline{\mathbf{w}}\| \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \mathbf{b}_{\ell, \underline{\mathbf{w}}}^G \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda_\ell \right)^{-1} \leq (8k \log(1 + \tau_{\ell-1}^G/\lambda))^2$. Thus we have taken

at most $\tau_\ell^G + \frac{8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})(8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) + 1)}{2}$ pulls. Thus, for any $\delta \in (0, 1)$ we have

$$\mathbb{P} \left(\bigcup_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \left\{ \left| \langle \underline{\mathbf{w}}, \hat{\theta}_\ell - \theta^* \rangle \right| \geq \sqrt{\frac{64B_*^\ell \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \log(2|\mathcal{W}|/\delta)}{\tau_\ell^G}} \right\} \right) \leq \delta.$$

The claim of the lemma follows. \square

Discussion 3. (Phase Length) It follows from Proposition 1 that if the gaps are known, one can set the phase length as

$$\tau_\ell^G = \frac{64B_*^\ell \rho(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) \log(2|\mathcal{W}|/\delta)}{(\underline{\mathbf{w}}^\top (\hat{\theta}_\ell - \theta^*))^2}$$

since $8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \leq \rho(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) := 8k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})$ to guarantee that the event $\bigcup_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \left\{ \left| \langle \underline{\mathbf{w}}, \hat{\theta}_\ell - \theta^* \rangle \right| \right\}$ holds with probability greater than $1 - \delta$.

However, since in practice, the gaps are not known, for an agnostic algorithm that does not know the gaps, one can set a proxy for the gap as ϵ_ℓ (for some $\epsilon_\ell > 0$) and get the phase length as follows:

$$\tau_\ell^G = \frac{64B_*^\ell \rho(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) \log(2|\mathcal{W}|/\delta)}{\epsilon_\ell^2}.$$

This gives us the desired phase length so that the event $\bigcup_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \left\{ \left| \langle \underline{\mathbf{w}}, \hat{\theta}_\ell - \theta^* \rangle \right| \right\}$ holds with probability greater than $1 - \delta$.

Lemma 10. Assume that $\max_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}} \langle \underline{\mathbf{w}}_* - \underline{\mathbf{w}}, \theta^* \rangle \leq 2$. With probability at least $1 - \delta$, we have $\underline{\mathbf{w}}_* \in \underline{\mathcal{W}}_\ell$ and $\max_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \langle \underline{\mathbf{w}}_* - \underline{\mathbf{w}}, \theta^* \rangle \leq 4\epsilon_\ell$ for all $\ell \in \mathbb{N}$.

Proof. For any $\mathcal{V} \subseteq \underline{\mathcal{W}}_\ell$ be the active set and $\underline{\mathbf{w}} \in \mathcal{V}$ define

$$\mathcal{E}_{\underline{\mathbf{w}}, \ell}(\mathcal{V}) = \left\{ \left| \langle \underline{\mathbf{w}} - \underline{\mathbf{w}}_*, \hat{\theta}_\ell(\mathcal{V}) - \theta^* \rangle \right| \leq \epsilon_\ell \right\} \quad (22)$$

where it is implicit that $\hat{\theta}_\ell := \hat{\theta}_\ell(\mathcal{V})$ is the design constructed in the algorithm at stage ℓ with respect to $\underline{\mathcal{W}}_\ell = \mathcal{V}$. Also note that $\delta_\ell = \frac{\delta}{4\ell^2}$. Given $\underline{\mathcal{W}}_\ell$, with probability at least $1 - 2 \cdot \frac{\delta}{4\ell^2 |\underline{\mathcal{W}}|}$

$$\begin{aligned}
\left| \langle \underline{\mathbf{w}} - \underline{\mathbf{w}}_*, \hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^* \rangle \right| &\stackrel{(a)}{\leq} \sqrt{64B_*^\ell k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}) \log(4\ell^2 |\underline{\mathcal{W}}|/\delta)} \\
&\leq \frac{(k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}))^2}{\sqrt{\tau_\ell^G}} \sqrt{64B_*^\ell \log(4\ell^2 |\underline{\mathcal{W}}|/\delta)} \\
&\stackrel{(b)}{\leq} \sqrt{\frac{\|\underline{\mathbf{w}} - \underline{\mathbf{w}}_*\|^2_{(\sum_{\underline{\mathbf{w}} \in \mathcal{V}} \mathbf{b}_{\ell, \underline{\mathbf{w}}}^G(\mathcal{V}) \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda_\ell)^{-1}}}{64B_*^\ell \epsilon_\ell^{-2} \rho(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) \log(4\ell^2 |\underline{\mathcal{W}}|/\delta)}} \sqrt{64B_*^\ell \log(4\ell^2 |\underline{\mathcal{W}}|/\delta)} \\
&\leq \sqrt{\frac{\|\underline{\mathbf{w}} - \underline{\mathbf{w}}_*\|^2_{(\sum_{\underline{\mathbf{w}} \in \mathcal{V}} \mathbf{b}_{\ell, \underline{\mathbf{w}}}^G(\mathcal{V}) \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda_\ell)^{-1}}}{\epsilon_\ell^{-2} \rho(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) \log(4\ell^2 |\underline{\mathcal{W}}|/\delta)}} \sqrt{\log(4\ell^2 |\underline{\mathcal{W}}|/\delta)} \\
&\stackrel{(c)}{=} \epsilon_\ell
\end{aligned}$$

where, (a) follows Proposition 1. The (b) follows as

$$(k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda}))^2 \leq (k \log(1 + \frac{\tau_\ell^G}{\lambda}))^2 := \rho(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) = \|\underline{\mathbf{w}} - \underline{\mathbf{w}}_*\|_{(\sum_{\underline{\mathbf{w}} \in \mathcal{V}} \mathbf{b}_{\ell, \underline{\mathbf{w}}}^G(\mathcal{V}) \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda_\ell)^{-1}}^2.$$

The (c) follows as $\rho(\mathcal{Y}(\underline{\mathcal{W}}_\ell)) = \|\underline{\mathbf{w}} - \underline{\mathbf{w}}_*\|_{(\sum_{\underline{\mathbf{w}} \in \mathcal{V}} \mathbf{b}_{\ell, \underline{\mathbf{w}}}^G(\mathcal{V}) \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda_\ell)^{-1}}^2$. By exactly the same sequence of steps as above, we have

$$\mathbb{P} \left(\bigcap_{\ell=1}^{\infty} \bigcap_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \left\{ \left| \langle \underline{\mathbf{w}} - \underline{\mathbf{w}}_*, \hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^* \rangle \right| > \epsilon_\ell \right\} \right) = \mathbb{P} \left(\bigcap_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell} \bigcap_{\ell=1}^{\infty} \mathcal{E}_{\underline{\mathbf{w}}, \ell}(\underline{\mathcal{W}}_\ell) \right) \geq 1 - \delta,$$

so assume these events hold. Consequently, for any $\underline{\mathbf{w}}' \in \underline{\mathcal{W}}_\ell$

$$\begin{aligned}
\langle \underline{\mathbf{w}}' - \underline{\mathbf{w}}_*, \hat{\boldsymbol{\theta}}_\ell \rangle &= \langle \underline{\mathbf{w}}' - \underline{\mathbf{w}}_*, \hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^* \rangle + \langle \underline{\mathbf{w}}' - \underline{\mathbf{w}}_*, \boldsymbol{\theta}^* \rangle \\
&\leq \langle \underline{\mathbf{w}}' - \underline{\mathbf{w}}_*, \hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^* \rangle \\
&\leq \epsilon_\ell
\end{aligned}$$

so that $\underline{\mathbf{w}}_*$ would survive to round $\ell + 1$. And for any $\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell$ such that $\langle \underline{\mathbf{w}}_* - \underline{\mathbf{w}}, \boldsymbol{\theta}^* \rangle > 2\epsilon_\ell$ we have

$$\begin{aligned}
\max_{\underline{\mathbf{w}}' \in \underline{\mathcal{W}}_\ell} \langle \underline{\mathbf{w}}' - \underline{\mathbf{w}}, \hat{\boldsymbol{\theta}}_\ell \rangle &\geq \langle \underline{\mathbf{w}}_* - \underline{\mathbf{w}}, \hat{\boldsymbol{\theta}}_\ell \rangle \\
&= \langle \underline{\mathbf{w}}_* - \underline{\mathbf{w}}, \hat{\boldsymbol{\theta}}_\ell - \boldsymbol{\theta}^* \rangle + \langle \underline{\mathbf{w}}_* - \underline{\mathbf{w}}, \boldsymbol{\theta}^* \rangle \\
&> -\epsilon_\ell + 2\epsilon_\ell \\
&= \epsilon_\ell
\end{aligned}$$

which implies this $\underline{\mathbf{w}}$ would be eliminated. Note that this implies that $\max_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_{\ell+1}} \langle \underline{\mathbf{w}}_* - \underline{\mathbf{w}}, \boldsymbol{\theta}^* \rangle \leq 2\epsilon_\ell = 4\epsilon_{\ell+1}$. Hence, the claim of the lemma follows. \square

A.5 Final Sample Complexity Bound for Single Task Setting

Theorem 1. (Restatement) With probability at least $1 - \delta$, **GOBLIN** returns the best arms $\mathbf{x}_*, \mathbf{z}_*$, and the number of samples used is bounded by

$$\tilde{O} \left(\frac{(d_1 + d_2)r}{\Delta^2} + \frac{\sqrt{d_1 d_2 r}}{S_r} \right)$$

where, $\Delta = \min_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}, \mathbf{z} \in \mathcal{Z} \setminus \{\mathbf{z}_*\}} (\mathbf{x}_*^\top \boldsymbol{\Theta}_* \mathbf{z}_* - \mathbf{x}^\top \boldsymbol{\Theta}_* \mathbf{z})$, and S_r is the r -th largest singular value of $\boldsymbol{\Theta}_*$.

Proof. For the rest of the proof we have that the good events $\mathcal{F}_\ell \cap \mathcal{E}_{\mathbf{w},\ell}(\mathcal{W}_\ell)$ holds true for each phase ℓ with probability greater than $(1 - \delta)$. The two events are defined in (16) and (22).

Second Stage: Define $\mathcal{A}_\ell = \{\mathbf{w} \in \mathcal{W}_\ell : \langle \mathbf{w}_* - \mathbf{w}, \boldsymbol{\theta}^* \rangle \leq 4\epsilon_\ell\}$. Note that by assumption $\mathcal{W} = \mathcal{W}_1 = \mathcal{S}_1$. The above lemma implies that with probability at least $1 - \delta$ we have $\bigcap_{\ell=1}^\infty \{\mathcal{W}_\ell \subseteq \mathcal{A}_\ell\}$. This implies that

$$\begin{aligned} \rho^G(\mathcal{W}_\ell) &= \min_{\mathbf{b} \in \Delta_{\mathcal{W}}} \max_{\mathbf{w}, \mathbf{w}' \in \mathcal{W}_\ell} \|\mathbf{w} - \mathbf{w}'\|_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}}^2 \\ &\leq \min_{\mathbf{b} \in \Delta_{\mathcal{W}}} \max_{\mathbf{w}, \mathbf{w}' \in \mathcal{A}_\ell} \|\mathbf{w} - \mathbf{w}'\|_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}}^2 \\ &= \rho^G(\mathcal{A}_\ell). \end{aligned}$$

Define $k_1^\ell = 8k \log(1 + \tau_{\ell-1}^G/\lambda)$. For $\ell \geq \lceil \log_2(4\Delta^{-1}) \rceil$ we have that $\mathcal{A}_\ell = \{\mathbf{w}_*\}$, thus, the sample complexity to identify \mathbf{w}_* is equal to

$$\begin{aligned} &\sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \sum_{\mathbf{w} \in \mathcal{W}} \left[\tau_\ell^G \widehat{\mathbf{b}}_{\ell, \mathbf{w}}^G \right] = \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{(k_1^\ell + 1)k_1^\ell}{2} + \tau_\ell^G \right) \\ &= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{(k_1^\ell + 1)k_1^\ell}{2} + 2\epsilon_\ell^{-2} \rho^G(\mathcal{W}_\ell) B_*^\ell \log(4k_1^\ell \ell^2 |\mathcal{W}|/\delta) \right) \\ &\stackrel{(a)}{\leq} 2 \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{(k+1)k}{2} \log^2(1 + \tau_{\ell-1}^G) + 2\epsilon_\ell^{-2} \rho^G(\mathcal{W}_\ell) B_*^\ell \log(4k_1^\ell \ell^2 |\mathcal{W}|/\delta) \right) \\ &\stackrel{(b)}{\leq} 2 \frac{(k+1)k}{2} \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} (\log^2(1 + \tau_{\ell-1}^G) + 8\epsilon_\ell^{-2} \rho^G(\mathcal{W}_\ell) B_*^\ell \log(4k_1^\ell \ell^2 |\mathcal{W}|/\delta)) \\ &\stackrel{(c)}{\leq} (k+1)k \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} (1 + 16\epsilon_\ell^{-2} \rho^G(\mathcal{W}_\ell) B_*^\ell \log^2(4k_1^\ell \ell^2 |\mathcal{W}|/\delta)) \\ &\stackrel{(d)}{\leq} (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 32\epsilon_\ell^{-2} f(\mathcal{A}_\ell) B_*^\ell \log(4k\ell^2 |\mathcal{W}|/\delta) \\ &\stackrel{(e)}{\leq} (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 32\epsilon_\ell^{-2} f(\mathcal{A}_\ell) (64\lambda S^2 + 64\tau_{\ell-1}^G) \log(4k\ell^2 |\mathcal{W}|/\delta) \\ &= (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 32\epsilon_\ell^{-2} f(\mathcal{A}_\ell) (64\lambda S^2) \log(4k\ell^2 |\mathcal{W}|/\delta) \\ &\quad + (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 32\epsilon_\ell^{-2} f(\mathcal{A}_\ell) (64\tau_{\ell-1}^G) \log(4k\ell^2 |\mathcal{W}|/\delta) \\ &\stackrel{(f)}{\leq} (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 64\epsilon_\ell^{-2} f(\mathcal{A}_\ell) (64\lambda S^2) \log(4k\ell^2 |\mathcal{W}|/\delta) \\ &\leq (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + 2048\lambda S^2 \log\left(\frac{4k \log_2^2(8\Delta^{-1}) |\mathcal{W}|}{\delta}\right) \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 2^{2\ell} f(\mathcal{A}_\ell) \end{aligned}$$

where, (a) follows as $\log^2(1 + \tau_{\ell-1}^G/\lambda) \leq \log^2(1 + \tau_{\ell-1}^G)$, (b) follows by noting that $\log(x \log(1 + x)) \leq 2 \log(x)$ for any $x > 1$. The (c) follows by subsuming the $\log^2(1 + \tau_{\ell-1}^G)$ into $2\tau_\ell^G$. The (d) follows as $\log(1 + \tau_{\ell-1}^G) < \tau_\ell^G$ which enables us to replace the k_1^ℓ inside the log with an additional

factor of 2. The (e) follows by noting that

$$\begin{aligned}
B_*^\ell &\leq 64(\sqrt{\lambda}S + \sqrt{\lambda_\ell^\perp}S_\ell^\perp) \\
&\leq 64\lambda S^2 + \left(\frac{64\tau_{\ell-1}^G}{8(d_1 + d_2)r \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})} \right) \cdot \left(\frac{8d_1d_2r}{\tau_\ell^E S_r^2} \log\left(\frac{d_1 + d_2}{\delta_\ell}\right) \right) \\
&\stackrel{(a_1)}{\leq} 64\lambda S^2 + 64\tau_{\ell-1}^G.
\end{aligned} \tag{23}$$

where, (a_1) follows by first substituting the value of $\tau_\ell^E := \frac{\sqrt{8d_1d_2r \log(4\ell^2|\mathcal{W}|/\delta_\ell)}}{S_r}$ and noting that $\sqrt{(d_1d_2r)} \leq (d_1 + d_2)r$ and cancelling out the other terms. Finally the (f) follows by subsuming the $\tau_{\ell-1}^G$ with a factor of 2 into the quantity of τ_ℓ^G . Then it follows that

$$\begin{aligned}
\rho_*^G &= \inf_{\mathbf{b} \in \Delta_{\mathcal{W}}} \max_{\mathbf{w} \in \mathcal{W}} \frac{\|\mathbf{w} - \mathbf{w}_*\|^2_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}}}{(\langle \mathbf{w} - \mathbf{w}_*, \boldsymbol{\theta}^* \rangle)^2} \\
&= \inf_{\mathbf{b} \in \Delta_{\mathcal{W}}} \max_{\ell \leq \lceil \log_2(4\Delta^{-1}) \rceil} \max_{\mathbf{w} \in \mathcal{A}_\ell} \frac{\|\mathbf{w} - \mathbf{w}_*\|^2_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}}}{(\langle \mathbf{w} - \mathbf{w}_*, \boldsymbol{\theta}^* \rangle)^2} \\
&\geq \frac{1}{\lceil \log_2(4\Delta^{-1}) \rceil} \inf_{\mathbf{b} \in \Delta_{\mathcal{W}}} \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \max_{\mathbf{w} \in \mathcal{A}_\ell} \frac{\|\mathbf{w} - \mathbf{w}_*\|^2_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}}}{(\langle \mathbf{w} - \mathbf{w}_*, \boldsymbol{\theta}^* \rangle)^2} \\
&\geq \frac{1}{16 \lceil \log_2(4\Delta^{-1}) \rceil} \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 2^{2\ell} \inf_{\mathbf{b} \in \Delta_{\mathcal{W}}} \max_{\mathbf{w} \in \mathcal{A}_\ell} \|\mathbf{w} - \mathbf{w}_*\|^2_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}} \\
&\geq \frac{1}{64 \lceil \log_2(4\Delta^{-1}) \rceil} \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 2^{2\ell} \inf_{\mathbf{b} \in \Delta_{\mathcal{W}}} \max_{\mathbf{w}, \mathbf{w}' \in \mathcal{A}_\ell} \|\mathbf{w} - \mathbf{w}'\|^2_{(\sum_{\mathbf{w} \in \mathcal{W}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}} \\
&\geq \frac{1}{64 \lceil \log_2(4\Delta^{-1}) \rceil} \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 2^{2\ell} f(\mathcal{A}_\ell).
\end{aligned}$$

This implies that

$$\sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 2^{2\ell} f(\mathcal{A}_\ell) \leq \rho_*^G 64 \lceil \log_2(4\Delta^{-1}) \rceil$$

Plugging this back we get

$$\begin{aligned}
\sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \sum_{\mathbf{w} \in \mathcal{W}} \left[\tau_\ell^G \widehat{\mathbf{b}}_{\ell, \mathbf{w}} \right] &\leq (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + 2048\lambda S^2 \log\left(\frac{8k \log_2^2(8\Delta^{-1}) |\mathcal{W}|}{\delta}\right) \rho_*^G 64 \lceil \log_2(4\Delta^{-1}) \rceil \\
&\leq (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + C_2 \lambda S^2 \log\left(\frac{8k \log_2^2(8\Delta^{-1}) |\mathcal{W}|}{\delta}\right) \rho_*^G \lceil \log_2(4\Delta^{-1}) \rceil
\end{aligned}$$

for some constant $C_2 > 0$. Now to understand the bound we need the following: Let $\text{conv}(\mathcal{W} \cup -\mathcal{W})$ denote the convex hull of $\mathcal{W} \cup -\mathcal{W}$, and for any set $\mathcal{Y} \subset \mathbb{R}^p$ define the gauge of \mathcal{Y}

$$\gamma_{\mathcal{Y}} = \max\{c > 0 : c\mathcal{Y} \subset \text{conv}(\mathcal{W} \cup -\mathcal{W})\} \tag{24}$$

In the case where \mathcal{Y} is a singleton $\mathcal{Y} = \{y\}$, $\gamma(y) := \gamma_{\mathcal{Y}}$ is the gauge norm of \mathbf{y} with respect to $\text{conv}(\underline{\mathcal{W}} \cup -\underline{\mathcal{W}})$. We can provide a natural upper bound for $\rho(\mathcal{Y})$ in terms of the gauge. Observe that

$$\begin{aligned} \rho_*^G &= \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}}} \max_{\mathbf{y} \in \mathcal{Y}} \|\mathbf{y}\|^2_{(\sum_{\mathbf{w} \in \underline{\mathcal{W}}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}} \\ &= \frac{1}{\gamma_{\mathcal{Y}}^2} \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}}} \max_{\mathbf{y} \in \mathcal{Y}} \|\mathbf{y} \gamma_{\mathcal{Y}}\|^2_{(\sum_{\mathbf{w} \in \underline{\mathcal{W}}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}} \\ &\leq \frac{1}{\gamma_{\mathcal{Y}}^2} \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}}} \max_{\mathbf{w} \in \text{conv}(\underline{\mathcal{W}} \cup -\underline{\mathcal{W}})} \|\mathbf{w}\|^2_{(\sum_{\mathbf{w} \in \underline{\mathcal{W}}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}} \\ &\stackrel{(a)}{=} \frac{1}{\gamma_{\mathcal{Y}}^2} \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}}} \max_{\mathbf{w} \in \underline{\mathcal{W}}} \|\mathbf{w}\|^2_{(\sum_{\mathbf{w} \in \underline{\mathcal{W}}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}} \stackrel{(b)}{\leq} k^{3/2} \\ &\leq \frac{k}{\gamma_{\mathcal{Y}}^2} O \left(B_* \log \left(\frac{k \log_2(\Delta^{-1}) |\mathcal{W}|}{\delta} \right) \lceil \log_2(\Delta^{-1}) \rceil \right) \end{aligned}$$

The (a) follows from the fact that the maximum value of a convex function on a convex set must occur at a vertex. The (b) follows from Kiefer-wolfowitz theorem for $\underline{\mathbf{w}} \in \mathbb{R}^p$ such that $\inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}}} \max_{\mathbf{w} \in \underline{\mathcal{W}}} \|\mathbf{w}\|^2_{(\sum_{\mathbf{w} \in \underline{\mathcal{W}}} \mathbf{b}_{\mathbf{w}} \mathbf{w} \mathbf{w}^\top + \Lambda)^{-1}} \leq k \log(1 + \frac{\tau_{\ell-1}^G}{\lambda})$. The simplified sample complexity for the second stage is given by

$$N_2 \leq O \left(\frac{k}{\Delta^2} \log \left(\frac{k \log_2(\Delta^{-1}) |\mathcal{W}|}{\delta} \right) \right) = \tilde{O} \left(\frac{(d_1 + d_2)r}{\Delta^2} \right)$$

where $\Delta = \min_{\mathbf{w} \in \underline{\mathcal{W}}} (\mathbf{w}_* - \mathbf{w})^\top \boldsymbol{\theta}_* \stackrel{(a_1)}{=} \min_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}, \mathbf{z} \in \mathcal{Z} \setminus \{\mathbf{z}_*\}} (\mathbf{x}_*^\top \boldsymbol{\Theta}_* \mathbf{z}_* - \mathbf{x}^\top \boldsymbol{\Theta}_* \mathbf{z})$. The (a₁) follows by reshaping the arms in $\underline{\mathcal{W}}$ to recover the arms in \mathcal{X} and \mathcal{Z} .

1st Stage: First recall that the E -optimal design in step 3 of Algorithm 1 satisfies the Assumption 4 as the sample distribution \mathcal{D} has finite second order moments. For the first stage first observe that by plugging in the definition of $\tau_\ell^E = \frac{\sqrt{8d_1 d_2 r \log(4\ell^2 |\mathcal{W}|/\delta)}}{S_r}$ we get

$$\begin{aligned} \|\boldsymbol{\theta}_{k+1:p}^*\|_2^2 &= \sum_{i>r \wedge j>r} H_{ij}^2 = \left\| (\widehat{\mathbf{U}}_\ell^\perp)^\top (\mathbf{U}^* \mathbf{S}^* \mathbf{V}^{*\top}) \widehat{\mathbf{V}}_\ell^\perp \right\|_F^2 \\ &\leq \left\| (\widehat{\mathbf{U}}_\ell^\perp)^\top \mathbf{U}^* \right\|_F^2 \|\mathbf{S}^*\|_2^2 \left\| (\widehat{\mathbf{V}}_\ell^\perp)^\top \mathbf{V}^* \right\|_F^2 \leq O \left(\frac{d_1 d_2 r}{\tau_\ell^E S_r^2} \log \left(\frac{d_1 + d_2}{\delta} \right) \right) \\ &= O \left(\frac{\sqrt{d_1 d_2 r}}{S_r} \log \left(\frac{d_1 + d_2}{\delta} \right) \right) \end{aligned}$$

which implies $\|\boldsymbol{\theta}_{k+1:p}^*\|_2 = \tilde{O}(\sqrt{d_1 d_2 r}/S_r)$. Now we bound the sample complexity from the first stage. From the first stage we can show that we have for the arm set $\overline{\mathcal{W}}$

$$\begin{aligned} N_1 &= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \sum_{\overline{\mathbf{w}} \in \overline{\mathcal{W}}} \left\lceil \tau_\ell^E \widehat{\mathbf{b}}_{\ell, \overline{\mathbf{w}}}^E \right\rceil \\ &= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{(p+1)p}{2} + \tau_\ell^E \right) \\ &= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{(p+1)p}{2} + \frac{\sqrt{8d_1 d_2 r \log(4\ell^2 |\mathcal{W}|/\delta)}}{S_r} \right) \\ &\leq (p+1)p \lceil \log_2(4\Delta^{-1}) \rceil + 32 \frac{\sqrt{d_1 d_2 r}}{S_r} \log \left(\frac{4 \log_2^2(8\Delta^{-1}) |\mathcal{W}|}{\delta} \right) \lceil \log_2(4\Delta^{-1}) \rceil \\ &\stackrel{(a)}{=} O \left(\frac{\sqrt{d_1 d_2 r}}{S_r} \log \left(\frac{4 \log_2^2(8\Delta^{-1}) |\mathcal{W}|}{\delta} \right) \right) = \tilde{O} \left(\frac{\sqrt{d_1 d_2 r}}{S_r} \right) \end{aligned}$$

where, (a) follows as $p = d_1 d_2$. Combining N_1 and N_2 gives the claim of the theorem. \square

A.6 Multi-Task Pure Exploration Proofs

Remark 2. (Comparison with Du et al. (2023)) In this remark, we discuss a key comparison of DouExpDes (Du et al., 2023) with GOBLIN. Note that DouExpDes does not implement the second stage of finding the $\mathbf{S}_{m,*} \in \mathbb{R}^{k_1 \times k_2}$ for each of the m bilinear bandits. Hence, DouExpDes does not rotate the arms so that the last $(k_1 - r) \cdot (k_2 - r)$ components are from the complementary subspaces of the left and right eigenvectors of $\mathbf{S}_{m,*}$. This results in DouExpDes suffering a sample complexity of $\tilde{O}(k_1 k_2 / \Delta^2)$ even though it learns the common feature extractors shared across the tasks. In contrast GOBLIN uses the second stage to learn $\mathbf{S}_{m,*} \in \mathbb{R}^{k_1 \times k_2}$ and reduces the latent bilinear bandits problem of $k_1 k_2$ dimension to $(k_1 + k_2)r$ dimension by rotating the arms so that the last $(k_1 - r) \cdot (k_2 - r)$ components are from the complementary subspaces of the left and right eigenvectors of $\mathbf{S}_{m,*}$. Hence, GOBLIN suffers a sample complexity of $\tilde{O}((k_1 + k_2)r / \Delta^2)$.

Remark 3. (Arm set) The observable left and right arm sets \mathcal{X} and \mathcal{Z} are common across the M tasks. This leads to each task estimating the same E-optimal design in line 3 of Algorithm 2 of stage 1. Note that Du et al. (2023) also uses a similar idea of the same arm set \mathcal{X} shared across tasks in the linear bandit setting. Observe that if each task has access to its own separate arm sets \mathcal{X}_m and \mathcal{Z}_m , then each of the m -tasks has to estimate a separate E-optimal design for the stage 1. This will lead to the sample complexity of the first stage scaling as $\tilde{O}(M\sqrt{d_1 d_2} / S_r)$ instead of $\tilde{O}(\sqrt{d_1 d_2} / S_r)$.

Good Event: We first recall the total stage 1 length as

$$\tau_\ell^E := \frac{\sqrt{8d_1 d_2 r \log(4\ell^2 |\mathcal{W}| / \delta_\ell)}}{S_r}.$$

Then define the good event \mathcal{F}_ℓ in phase ℓ that GOBLIN has a good estimate of $\mathbf{Z}_* = \frac{1}{M} \sum_{m=1}^M \boldsymbol{\Theta}_m$ as follows: For any phase $\ell > 0$

$$\mathcal{F}_\ell := \left\{ \left\| \hat{\mathbf{Z}}_\ell - \mu^* \mathbf{Z}_* \right\|_F^2 \leq \frac{C_1 d_1 d_2 r \log\left(\frac{2(d_1 + d_2)}{\delta_\ell}\right)}{\tau_\ell^E} \right\}, \quad (25)$$

where, $C_1 = 36(4 + S_0^2)C$, $\|\mathbf{X}\|_F, \|\boldsymbol{\Theta}_*\|_F \leq S_0$, some nonzero constant μ^* , $\mathbb{E}[(\mathcal{S}^P(\mathbf{X}))_{ij}^2] \leq C, \forall i, j$, and $\hat{\boldsymbol{\Theta}}_\ell$ is the estimate from (8). Then define the event

$$\mathcal{F} := \bigcap_{\ell=1}^{\infty} \mathcal{F}_\ell \quad (26)$$

Then we start by modifying Lemma 3 for the multi-task setting. We first prove this support lemma for the loss function defined in (8).

Lemma 11. Let $L : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}$ is the loss function defined in (8). Then by setting

$$t = \sqrt{2d_1 d_2 C(4 + S_0^2) \log\left(\frac{2(d_1 + d_2)}{\delta_\ell}\right)},$$

$$\nu = \frac{t}{(4 + S_0) C d_1 d_2 \sqrt{M \tau_\ell^E}} = \sqrt{\frac{2 \log\left(\frac{2(d_1 + d_2)}{\delta_\ell}\right)}{M \tau_\ell^E d_1 d_2 C(4 + S_0^2)}},$$

we have with probability at least $1 - \delta_\ell$, it holds that

$$\mathbb{P} \left(\left\| \nabla L(\mu^* \mathbf{Z}_*) \right\|_{op} \geq \frac{2t}{\sqrt{M \tau_\ell^E}} \right) \leq \delta_\ell,$$

where $\mu^* = \frac{2}{M} \mathbb{E}[\langle \mathbf{X}_m, \mathbf{Z}_* \rangle] > 0$, and $\mathbf{X}_m = \mathbf{x}_m \mathbf{z}_m^\top$.

Proof. Let $\mathbf{Z}_* = \frac{1}{M} \sum_{m=1}^M \boldsymbol{\Theta}_{m,*}$. Let $\mathbf{X}_{m,i} = \mathbf{x}_{m,i} \mathbf{z}_{m,i}^\top$ for $i \in [\tau_\ell^E]$. Based on the definition of our loss function $L(\cdot)$ in (8), we have that

$$\begin{aligned} \nabla_{x_m} L(\mathbf{Z}_*) &= \mu^* \mathbf{Z}_* - \frac{2}{M \tau_\ell^E} \sum_{m=1}^M \sum_{i=1}^{\tau_\ell^E} \tilde{\psi}_\nu(r_{m,i} \cdot Q(x_m)) \\ &= \frac{2}{M} \mathbb{E}[\langle \mathbf{X}_{m,1}, \mathbf{Z}_* \rangle] \mathbf{Z}_* - \frac{2}{M \tau_\ell^E} \sum_{m=1}^M \sum_{i=1}^{\tau_\ell^E} \tilde{\psi}_\nu(r_{m,i} \cdot Q(\mathbf{X}_{m,i})) \\ &\stackrel{(a)}{=} \frac{2}{M} \mathbb{E}[\langle \mathbf{X}_{m,1}, \mathbf{Z}_* \rangle Q(\mathbf{X}_{m,1})] - \frac{2}{M \tau_\ell^E} \sum_{m=1}^M \sum_{i=1}^{\tau_\ell^E} \tilde{\psi}_\nu(r_{m,i} \cdot Q(\mathbf{X}_{m,i})) \\ &\stackrel{(b)}{=} \frac{2}{M} \left[\mathbb{E}(r_{m,1} \cdot Q(\mathbf{X}_{m,1})) - \frac{1}{M \tau_\ell^E} \sum_{m=1}^M \sum_{i=1}^{\tau_\ell^E} \tilde{\psi}_\nu(r_{m,i} \cdot Q(\mathbf{X}_{m,i})) \right] \end{aligned}$$

where we have (a) due to the generalized Stein's Lemma stated in Lemma 1, and (b) comes from the fact that the random noise $\eta_1 = y_1 - \langle \mathbf{X}_1, \mathbf{Z}_* \rangle$ is zero-mean and independent from \mathbf{X}_1 . Therefore, in order to implement the Lemma 2, we can see that it suffices to get σ^2 defined as:

$$\begin{aligned} \sigma^2 &= \max \left(\left\| \frac{2}{M} \sum_{m=1}^M \sum_{j=1}^{\tau_\ell^E} \mathbb{E} \left[r_{m,j}^2 Q(\mathbf{X}_{m,j}) Q(\mathbf{X}_{m,j})^\top \right] \right\|_{\text{op}}, \left\| \frac{2}{M} \sum_{m=1}^M \sum_{j=1}^{\tau_\ell^E} \mathbb{E} \left[r_{m,j}^2 Q(\mathbf{X}_{m,j})^\top Q(\mathbf{X}_{m,j}) \right] \right\|_{\text{op}} \right) \\ &\leq \left\| \frac{2}{M} \sum_{m=1}^M \sum_{j=1}^{\tau_\ell^E} \mathbb{E} \left[r_{m,j}^2 Q(\mathbf{X}_{m,j}) Q(\mathbf{X}_{m,j})^\top \right] \right\|_{\text{op}} \leq \tau_\ell^E \times \left\| \mathbb{E} \left[r_{m,1}^2 Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{op}} \\ &\stackrel{(a)}{=} \tau_\ell^E \times \left\| \mathbb{E} \left[(\eta_{m,1} + \langle \mathbf{X}_{m,1}, \mathbf{Z}_* \rangle)^2 Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{op}} \\ &\stackrel{(b)}{=} \tau_\ell^E \times \left\| \mathbb{E} \left[\eta_{m,1}^2 Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] + \mathbb{E} \left[\langle \mathbf{X}_{m,1}, \mathbf{Z}_* \rangle^2 Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{op}} \\ &\stackrel{(c)}{=} \tau_\ell^E \times \left\| \mathbb{E}(\eta_{m,1}^2) \mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] + \mathbb{E} \left[\langle \mathbf{X}_{m,1}, \mathbf{Z}_* \rangle^2 Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{op}} \\ &\stackrel{(d)}{\leq} \tau_\ell^E \times \left\| 4 \mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] + S_0^2 \mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{op}} \\ &= (4 + S_0^2) \tau_\ell^E \times \left\| \mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{op}} \end{aligned}$$

where the (a) follows by plugging in the definition for reward, (b) follows by the linearity of expectation, (c) follows as noises are independent, and the inequality (d) comes from the fact that $|\langle \mathbf{X}_{m,1}, \mathbf{Z}_* \rangle| \leq S_0$, and $Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top$ is always positive semidefinite. Next, since we know that $\mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right]$ is always symmetric and positive semidefinite, and hence we have

$$\begin{aligned} \left\| \mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{op}} &\stackrel{(a)}{\leq} \left\| \mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right\|_{\text{nuc}} = \text{trace} \left(\mathbb{E} \left[Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right] \right) \\ &= \mathbb{E} \left[\text{trace} \left(Q(\mathbf{X}_{m,1}) Q(\mathbf{X}_{m,1})^\top \right) \right] = \mathbb{E} \left(\frac{2}{M} \sum_{m=1}^M \sum_{i=1}^{d_1} \sum_{j=1}^{d_2} Q_{ij}(\mathbf{X}_{m,1})^2 \right) \\ &\leq d_1 d_2 C. \end{aligned}$$

where, in (a) $\|\cdot\|_{\text{nuc}}$ denotes the nuclear norm. Therefore, we have that under 1-subGaussian assumption

$$\left\| \frac{2}{M} \sum_{m=1}^M \sum_{j=1}^{\tau_\ell^E} \mathbb{E} \left[r_{m,j}^2 Q(\mathbf{X}_{m,j}) Q(\mathbf{X}_{m,j})^\top \right] \right\|_{\text{op}} \leq (4 + S_0^2) d_1 d_2 \tau_\ell^E C.$$

And similarly, we can prove that

$$\left\| \frac{2}{M} \sum_{m=1}^M \sum_{j=1}^{\tau_\ell^E} \mathbb{E} \left[r_{m,j}^2 Q(\mathbf{X}_{m,j})^\top Q(\mathbf{X}_{m,j}) \right] \right\|_{\text{op}} \leq (4 + S_0^2) d_1 d_2 \tau_\ell^E C.$$

Therefore, we can take $\sigma^2 = (4 + S_0^2) d_1 d_2 \tau_\ell^E C$ consequently. By using Lemma 2, we have

$$\mathbb{P} \left(\|\nabla L(\mu^* \mathbf{Z}_*)\|_{\text{op}} \geq \frac{2t}{\sqrt{M \tau_\ell^E}} \right) \leq 2(d_1 + d_2) \exp \left(-\nu t \sqrt{M \tau_\ell^E} + \frac{\nu^2 (4 + S_0^2) C d_1 d_2 \tau_\ell^E}{2} \right)$$

By plugging the values of t and ν in Lemma 11, we finish the proof. \square

Lemma 12. *For any low-rank linear model with samples $\mathbf{X}_1 \dots, \mathbf{X}_{\tau_\ell^E}$ drawn from \mathcal{X} according to \mathcal{D} then for the optimal solution to the nuclear norm regularization problem in (2) with $\nu = \sqrt{2 \log(2(d_1 + d_2)/\delta_\ell) / ((4 + S_0^2) M \tau_\ell^E d_1 d_2)}$ and*

$$\gamma_\ell = 4 \sqrt{\frac{2(4 + S_0^2) C d_1 d_2 \log(2(d_1 + d_2)/\delta_\ell)}{M \tau_\ell^E}},$$

with probability at least $1 - \delta_\ell$ it holds that:

$$\left\| \hat{\mathbf{Z}}_\ell - \mu^* \mathbf{Z}_* \right\|_F^2 \leq \frac{C_1 d_1 d_2 r \log\left(\frac{2(d_1 + d_2)}{\delta_\ell}\right)}{M \tau_\ell^E},$$

for $C_1 = 36(4 + S_0^2)C$, $\|\mathbf{X}\|_F, \|\mathbf{Z}_*\|_F \leq S_0$, some nonzero constant μ^* , and $\mathbb{E} \left[(S^{\mathbf{p}}(\mathbf{X}))_{ij}^2 \right] \leq C, \forall i, j$. Summing over all phases $\ell \geq 1$ it follows that $\mathbb{P}(\mathcal{F}) \geq 1 - \delta/2$.

Proof. Since the estimator $\hat{\mathbf{Z}}_\ell$ minimizes the regularized loss function defined in (8), we have

$$L(\hat{\mathbf{Z}}_\ell) + \gamma_\ell \|\hat{\mathbf{Z}}_\ell\|_{\text{nuc}} \leq L(\mu^* \mathbf{Z}_*) + \gamma_\ell \|\mu^* \mathbf{Z}_*\|_{\text{nuc}}$$

And due to the fact that $L(\cdot)$ is a quadratic function, we have the following expression based on multivariate Taylor's expansion:

$$L(\hat{\mathbf{Z}}_\ell) - L(\mu^* \mathbf{Z}_*) = \langle \nabla L(\mu^* \mathbf{Z}_*), \Theta \rangle + 2\|\Theta\|_F^2, \quad \text{where } \Theta = \hat{\mathbf{Z}}_\ell - \mu^* \mathbf{Z}_*$$

By rearranging the above two results, we can deduce that

$$\begin{aligned} 2\|\Theta\|_F^2 &\leq -\langle \nabla L(\mu^* \mathbf{Z}_*), \Theta \rangle + \gamma_\ell \|\mu^* \mathbf{Z}_*\|_{\text{nuc}} - \gamma_\ell \|\hat{\mathbf{Z}}_\ell\|_{\text{nuc}} \\ &\stackrel{(a)}{\leq} \|\nabla L(\mu^* \mathbf{Z}_*)\|_{\text{op}} \|\Theta\|_{\text{nuc}} + \gamma_\ell \|\mu^* \mathbf{Z}_*\|_{\text{nuc}} - \gamma_\ell \|\hat{\mathbf{Z}}_\ell\|_{\text{nuc}}, \end{aligned} \quad (27)$$

where (a) comes from the duality between matrix operator norm and nuclear norm. Next, we represent the saturated SVD of \mathbf{Z}_* in the main paper as $\mathbf{Z}_* = \mathbf{U} \mathbf{D} \mathbf{V}^\top$ where $\mathbf{U} \in \mathbb{R}^{d_1 \times r}$ and $\mathbf{V} \in \mathbb{R}^{d_2 \times r}$, and here we would work on its full version, i.e.

$$\mathbf{Z}_* = (\mathbf{U}, \mathbf{U}_\perp) \begin{pmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} (\mathbf{V}, \mathbf{V}_\perp)^\top = (\mathbf{U}, \mathbf{U}_\perp) \mathbf{D}^* (\mathbf{V}, \mathbf{V}_\perp)^\top$$

where we have $\mathbf{U}_\perp \in \mathbb{R}^{d_1 \times (d_1 - r)}$, $\mathbf{D}^* \in \mathbb{R}^{d_1 \times d_2}$ and $\mathbf{V}_\perp \in \mathbb{R}^{d_2 \times (d_2 - r)}$. Furthermore, we define

$$\Lambda = (\mathbf{U}, \mathbf{U}_\perp)^\top \Theta (\mathbf{V}, \mathbf{V}_\perp) = \begin{pmatrix} \mathbf{U}^\top \Theta \mathbf{V} & \mathbf{U}^\top \Theta \mathbf{V}_\perp \\ \mathbf{U}_\perp^\top \Theta \mathbf{V} & \mathbf{U}_\perp^\top \Theta \mathbf{V}_\perp \end{pmatrix} = \Lambda_1 + \Lambda_2$$

where we write

$$\Lambda_1 = \begin{pmatrix} 0 & 0 \\ 0 & \mathbf{U}_\perp^\top \Theta \mathbf{V}_\perp \end{pmatrix}, \quad \Lambda_2 = \begin{pmatrix} \mathbf{U}^\top \Theta \mathbf{V} & \mathbf{U}^\top \Theta \mathbf{V}_\perp \\ \mathbf{U}_\perp^\top \Theta \mathbf{V} & 0 \end{pmatrix}.$$

Afterward, it holds that

$$\begin{aligned}
\|\widehat{\mathbf{Z}}_\ell\|_{\text{nuc}} &= \|\mu^* \mathbf{Z}_* + \boldsymbol{\Theta}\|_{\text{nuc}} \stackrel{(a)}{=} \left\| (\mathbf{U}, \mathbf{U}_\perp) (\mu^* \mathbf{D}^* + \boldsymbol{\Lambda}) (\mathbf{V}, \mathbf{V}_\perp)^\top \right\|_{\text{nuc}} \\
&\stackrel{(b)}{=} \|\mu^* \mathbf{D}^* + \boldsymbol{\Lambda}\|_{\text{nuc}} + \|\mu^* \mathbf{D}^* + \boldsymbol{\Lambda}_1 + \boldsymbol{\Lambda}_2\|_{\text{nuc}} \\
&\geq \|\mu^* \mathbf{D}^* + \boldsymbol{\Lambda}_1\|_{\text{nuc}} - \|\boldsymbol{\Lambda}_2\|_{\text{nuc}} \\
&= \|\mu^* \mathbf{D}\|_{\text{nuc}} + \|\boldsymbol{\Lambda}_1\|_{\text{nuc}} - \|\boldsymbol{\Lambda}_2\|_{\text{nuc}} \\
&= \|\mu^* \mathbf{Z}_*\|_{\text{nuc}} + \|\boldsymbol{\Lambda}_1\|_{\text{nuc}} - \|\boldsymbol{\Lambda}_2\|_{\text{nuc}}, \tag{28}
\end{aligned}$$

where, (a) follows from the definition of \mathbf{Z}_* , and (b) follows from the definition of $\boldsymbol{\Lambda}$. This implies that

$$\|\mu^* \mathbf{Z}_*\|_{\text{nuc}} - \|\widehat{\mathbf{Z}}_\ell\|_{\text{nuc}} \leq \|\boldsymbol{\Lambda}_2\|_{\text{nuc}} - \|\boldsymbol{\Lambda}_1\|_{\text{nuc}}.$$

Combining (27) and (28), we have that

$$2\|\boldsymbol{\Theta}\|_F^2 \leq \left(\|\nabla L(\mu^* \mathbf{Z}_*)\|_{\text{op}} + \gamma_\ell \right) \|\boldsymbol{\Lambda}_2\|_{\text{nuc}} + \left(\|\nabla L(\mu^* \mathbf{Z}_*)\|_{\text{op}} - \gamma_\ell \right) \|\boldsymbol{\Lambda}_1\|_{\text{nuc}}.$$

Then, we refer to the setting in our Lemma 11, and we choose $\gamma_\ell = 4t/\sqrt{M\tau_\ell^E}$ where the value of t is determined in Lemma 11, i.e.

$$\gamma_\ell = 4\sqrt{\frac{2(4 + S_0^2)Cd_1d_2 \log(2(d_1 + d_2)/\delta_\ell)}{M\tau_\ell^E}},$$

we know that $\lambda_{T-1} \geq 2\|\nabla L(\mu^* \mathbf{Z}_*)\|_{\text{op}}$ with probability at least $1 - \delta_\ell$ for any $\delta_\ell \in (0, 1)$. Therefore, with a probability at least $1 - \delta_\ell$, we have

$$2\|\boldsymbol{\Theta}\|_F^2 \leq \frac{3}{2}\gamma_\ell \|\boldsymbol{\Lambda}_2\|_{\text{nuc}} - \frac{1}{2}\gamma_\ell \|\boldsymbol{\Lambda}_1\|_{\text{nuc}} \leq \frac{3}{2}\gamma_\ell \|\boldsymbol{\Lambda}_2\|_{\text{nuc}}.$$

Since we can easily verify that the rank of $\boldsymbol{\Lambda}_2$ is at most $2r$, and by using Cauchy-Schwarz Inequality we have that

$$2\|\boldsymbol{\Theta}\|_F^2 \leq \frac{3}{2}\gamma_\ell \sqrt{2r} \|\boldsymbol{\Lambda}_2\|_F \leq \frac{3}{2}\gamma_\ell \sqrt{2r} \|\boldsymbol{\Lambda}\|_F = \frac{3}{2}\gamma_\ell \sqrt{2r} \|\boldsymbol{\Theta}\|_F$$

which implies that

$$\|\boldsymbol{\Theta}\|_F \leq \frac{3}{4}\sqrt{2r}\gamma_\ell = 6\sqrt{\frac{(4 + S_0^2)Cd_1d_2r \log\left(\frac{2(d_1+d_2)}{\delta_\ell}\right)}{M\tau_\ell^E}}.$$

This implies that $\mathbb{P}(\mathcal{F}_\ell) \geq 1 - \delta_\ell$. Taking a union bound over all phases $\ell \geq 1$ and recalling $\delta_\ell := \frac{\delta}{2\ell^2}$, we obtain

$$\begin{aligned}
\mathbb{P}(\mathcal{F}) &\geq 1 - \sum_{\ell=1}^{\infty} \mathbb{P}(\mathcal{F}_\ell^c) \\
&\geq 1 - \sum_{\ell=1}^{\infty} \frac{\delta_\ell}{2} \\
&= 1 - \sum_{\ell=1}^{\infty} \frac{\delta}{4\ell^2} \\
&\geq 1 - \frac{\delta}{2}.
\end{aligned}$$

This concludes our proof. \square

Define $\mathbf{X}_{\text{batch}}^+ := (\mathbf{X}_{\text{batch}}^\top \mathbf{X}_{\text{batch}})^{-1} \mathbf{X}_{\text{batch}}^\top$ where $\mathbf{X}_{\text{batch}}^+$ is constructed through the E -optimal design. Using Lemma C.1 from Du et al. (2023) it holds that

$$\|\mathbf{X}_{\text{batch}}^+\| \leq \sqrt{\frac{(1 + \beta)\rho_\ell^E}{\bar{p}}}.$$

where $\bar{p} = 180d_1d_2/\beta^2$ is the batch size to control the rounding procedure and $\rho_\ell^E = \min_{\mathbf{b} \in \Delta_{\overline{\mathcal{W}}}} \|(\sum_{\overline{\mathbf{w}} \in \overline{\mathcal{W}}} \mathbf{b}_{\overline{\mathbf{w}}} \overline{\mathbf{w}} \overline{\mathbf{w}}^\top)^{-1}\|$. It follows then that $\|\mathbf{X}_{\text{batch}}^+\|^2 \leq 4\rho_\ell^E$.

Lemma 13. (Expectation of $\widehat{\mathbf{Z}}_\ell$). It holds that $\mathbb{E} [\widehat{\mathbf{Z}}_\ell] = \mathbf{Z} = \frac{1}{M} \sum_{m=1}^M \boldsymbol{\Theta}_m$.

Proof. Note that we can re-write

$$\widehat{\mathbf{Z}}_\ell = \arg \min_{\boldsymbol{\Theta} \in \mathbb{R}^{d_1 \times d_2}} L_\ell(\boldsymbol{\Theta}) + \gamma_\ell \|\boldsymbol{\Theta}\|_{\text{nuc}}, L_\ell(\boldsymbol{\Theta}) = \langle \boldsymbol{\Theta}, \boldsymbol{\Theta} \rangle - \frac{2}{M\tau_\ell^E} \sum_{m=1}^M \sum_{s=1}^{\tau_\ell^E} \langle \tilde{\psi}_\nu(r_{m,s} \cdot Q(\mathbf{x}_{m,s} \mathbf{z}_{m,s}^\top)), \boldsymbol{\Theta} \rangle$$

such that

$$\widehat{\mathbf{Z}}_\ell = \frac{2}{M\tau_\ell^E} \sum_{m=1}^M \sum_{s=1}^{\tau_\ell^E} \widehat{\boldsymbol{\Theta}}_{m,s} - \mathbf{X}_{\text{batch}}^+ (\mathbf{X}_{\text{batch}}^+)^{\top}$$

where $\widehat{\boldsymbol{\Theta}}_{m,s} = \langle \boldsymbol{\Theta}, \boldsymbol{\Theta} \rangle - \frac{2}{Ms} \sum_{m=1}^M \langle \tilde{\psi}_\nu(r_{m,s} \cdot Q(\mathbf{x}_{m,s} \mathbf{z}_{m,s}^\top)), \boldsymbol{\Theta} \rangle$. Now using Lemma C.2 from [Du et al. \(2023\)](#) we can prove the result of the lemma. \square

Lemma 14. (Concentration of $\widehat{\mathbf{B}}_{1,\ell}$). Suppose that event \mathcal{F}_ℓ holds. Then, for any phase $\ell > 0$,

$$\left\| (\widehat{\mathbf{B}}_{1,\ell}^\perp)^\top \mathbf{B}_1 \right\| \leq \frac{c' \rho_\ell^E \sqrt{(d_1 + d_2)r}}{S_r \sqrt{M\tau_\ell^E}} \log \left(\frac{16(d_1 + d_2)rM\tau_\ell^E}{\delta_\ell} \right),$$

for some constant $c' > 0$ and $\rho_\ell^E = \min_{\mathbf{b} \in \Delta_{\overline{\mathbf{W}}} } \|(\sum_{\overline{\mathbf{w}} \in \overline{\mathbf{W}}} \mathbf{b}_{\overline{\mathbf{w}}} \overline{\mathbf{w}} \overline{\mathbf{w}}^\top)^{-1}\|$.

Proof. Using the Davis-Kahan sin θ Theorem ([Bhatia, 2013](#)) and letting τ_ℓ^E be large enough to satisfy

$$\left\| \widehat{\mathbf{Z}}_\ell - \mu^* \mathbf{Z}_* \right\|_F^2 \leq \frac{C_1 d_1 d_2 r \log \left(\frac{2(d_1 + d_2)}{\delta_\ell} \right)}{M\tau_\ell^E}, \text{ we have}$$

$$\begin{aligned} \left\| (\widehat{\mathbf{B}}_{1,\ell}^\perp)^\top \mathbf{B}_1 \right\| &\leq \frac{\left\| \widehat{\mathbf{Z}}_\ell - \mathbb{E} [\widehat{\mathbf{Z}}_\ell] \right\|}{\sigma_r \left(\mathbb{E} [\widehat{\mathbf{Z}}_\ell] \right) - \sigma_{r+1} \left(\mathbb{E} [\widehat{\mathbf{Z}}_\ell] \right) - \left\| \widehat{\mathbf{Z}}_\ell - \mathbb{E} [\widehat{\mathbf{Z}}_\ell] \right\|} \\ &\stackrel{(a)}{\leq} \frac{c_0}{S_r} \left\| \widehat{\mathbf{Z}}_\ell - \mathbb{E} [\widehat{\mathbf{Z}}_\ell] \right\| \\ &\stackrel{(b)}{\leq} \frac{cc_0 \left\| \mathbf{X}_{\text{batch}}^+ \right\|^2 \sqrt{(d_1 + d_2)r}}{S_r \sqrt{M\tau_\ell^E}} \log \left(\frac{16(d_1 + d_2)rM\tau_\ell^E}{\delta_\ell} \right) \\ &\stackrel{(c)}{\leq} \frac{c' \rho_\ell^E \sqrt{(d_1 + d_2)r}}{S_r \sqrt{M\tau_\ell^E}} \log \left(\frac{16(d_1 + d_2)rM\tau_\ell^E}{\delta_\ell} \right). \end{aligned}$$

where, (a) follows from Assumption 2, the (b) follows from event \mathcal{F}_ℓ and (c) follows as $\left\| \mathbf{X}_{\text{batch}}^+ \right\|^2 \leq 4\rho_\ell^E$. The claim of the lemma follows. \square

Lemma 15. (Concentration of $\widehat{\mathbf{B}}_{2,\ell}$). Suppose that event \mathcal{F}_ℓ holds. Then, for any phase $\ell > 0$,

$$\left\| (\widehat{\mathbf{B}}_{2,\ell}^\perp)^\top \mathbf{B}_2 \right\| \leq \frac{c\rho_\ell^E \sqrt{(d_1 + d_2)r}}{S_r \sqrt{M\tau_\ell^E}} \log \left(\frac{16(d_1 + d_2)rM\tau_\ell^E}{\delta_\ell} \right),$$

for some constant $c' > 0$ and $\rho_\ell^E = \min_{\mathbf{b} \in \Delta_{\overline{\mathbf{W}}} } \|(\sum_{\overline{\mathbf{w}} \in \overline{\mathbf{W}}} \mathbf{b}_{\overline{\mathbf{w}}} \overline{\mathbf{w}} \overline{\mathbf{w}}^\top)^{-1}\|$.

Proof. The proof follows the same way as Lemma 14 and using the Davis-Kahan sin θ Theorem ([Bhatia, 2013](#))

$$\left\| (\widehat{\mathbf{B}}_{2,\ell}^\perp)^\top \mathbf{B}_2 \right\| \leq \frac{c\rho_\ell^E \sqrt{(d_1 + d_2)r}}{S_r \sqrt{M\tau_\ell^E}} \log \left(\frac{16(d_1 + d_2)rM\tau_\ell^E}{\delta_\ell} \right).$$

The claim of the lemma follows. \square

Good Event per Task: We now define the good event \mathcal{F}'_ℓ in phase ℓ that **GOBLIN** has a good estimate of $\mathbf{S}_{m,*}$ as follows: For any phase $\ell > 0$

$$\mathcal{F}'_\ell := \left\{ \left\| \hat{\mathbf{S}}_{m,\ell} - \mu^* \mathbf{S}_{m,*} \right\|_F^2 \leq \frac{C_1 k_1 k_2 r \log \left(\frac{2(k_1+k_2)}{\delta_\ell} \right)}{\tau_{m,\ell}^E} \right\}, \quad (29)$$

where, $C, \mu^* > 0$ are constants and $\hat{\mathbf{S}}_{m,\ell}$ is the estimate from (9). Then define the event

$$\mathcal{F}' := \bigcap_{\ell=1}^{\infty} \mathcal{F}'_\ell. \quad (30)$$

We now prove the following lemmas to show the good event \mathcal{F}'_ℓ holds with probability $(1 - \delta_\ell)$. Before proving the concentration of $\hat{\mathbf{S}}_{m,\ell}$ we first need to show that $\sigma_{\min}(\sum_{\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}} \mathbf{b}_{\tilde{\mathbf{w}}} \tilde{\mathbf{w}} \tilde{\mathbf{w}}^\top) > 0$. If this holds true then we can sample following E -optimal design.

Lemma 16. *For any phase $\ell > 0$ and task $m \in [M]$, let $\left\| \hat{\mathbf{B}}_{1,\ell}^\top \mathbf{B}_1^\perp \right\| \leq c_1$ and $\left\| \hat{\mathbf{B}}_{2,\ell}^\top \mathbf{B}_2^\perp \right\| \leq c_2$, for some $c_1, c_2 > 0$. Then we have*

$$\sigma_{\min} \left(\sum_{\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}} \mathbf{b}_{\tilde{\mathbf{w}}} \tilde{\mathbf{w}} \tilde{\mathbf{w}}^\top \right) > 0$$

Proof. We can show that

$$\sum_{\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}} \mathbf{b}_{\tilde{\mathbf{w}}} \tilde{\mathbf{w}} \tilde{\mathbf{w}}^\top \stackrel{(a)}{=} \sum_{\mathbf{x} \in \mathcal{X}_m, \mathbf{z} \in \mathcal{Z}_m} \mathbf{b}_{x,z} \underbrace{\hat{\mathbf{U}}_\ell^\top \mathbf{x}_m}_{\tilde{\mathbf{g}}_m} \underbrace{\mathbf{z}_m \hat{\mathbf{V}}_\ell^\top}_{\tilde{\mathbf{v}}_m^\top}$$

where, in (a) the $\mathbf{b}_{x,z}$ is the sampling proportion for the arms \mathbf{x} and \mathbf{z} (they are allocated the same proportion, as they are pulled the same number of times). Also note that from Lemma 12 we know that $\left\| \hat{\mathbf{B}}_{1,\ell}^\top \mathbf{B}_1^\perp \right\| \leq c_1$ and $\left\| \hat{\mathbf{B}}_{2,\ell}^\top \mathbf{B}_2^\perp \right\| \leq c_2$ for some $c_1, c_2 > 0$ holds with high probability. This helps us to apply Lemma 17 to get the claim of the lemma. \square

Lemma 17. (Restatement of Lemma C.5 from Du et al. (2023)) *For any phase $\ell > 0$ and task $m \in [M]$, if $\left\| \hat{\mathbf{U}}_\ell^\top \mathbf{U}^\perp \right\| \leq c$ for some $c > 0$, then we have*

$$\sigma_{\min} \left(\sum_{i=1}^n \mathbf{b}_m^*(\mathbf{x}_i) \hat{\mathbf{U}}_\ell^\top \mathbf{x}_i \mathbf{x}_i^\top \hat{\mathbf{U}}_\ell \right) > 0$$

where \mathbf{b}_m^* is a sampling proportion on \mathbf{x} .

Lemma 18. *Let $L' : \mathbb{R}^{k_1 \times k_2} \rightarrow \mathbb{R}$ is the loss function defined in (9). Then by setting*

$$t = \sqrt{2k_1 k_2 C (4 + S_0^2) \log \left(\frac{2(k_1 + k_2)}{\delta_\ell} \right)},$$

$$\nu = \frac{t}{(4 + S_0) C k_1 k_2 \sqrt{\tau_\ell^E}} = \sqrt{\frac{2 \log \left(\frac{2(k_1 + k_2)}{\delta_\ell} \right)}{\tau_\ell^E k_1 k_2 C (4 + S_0^2)}},$$

we have with probability at least $1 - \delta_\ell$, it holds that

$$\mathbb{P} \left(\left\| \nabla L'(\mu^* \mathbf{S}_{m,*}) \right\|_{op} \geq \frac{2t}{\sqrt{\tau_\ell^E}} \right) \leq \delta_\ell,$$

where $\mu^* = \mathbb{E}[\langle \mathbf{X}_m, \mathbf{S}_{m,*} \rangle] > 0$, and $\mathbf{X}_m = \tilde{\mathbf{g}}_m \tilde{\mathbf{v}}_m^\top$.

Proof. Let $\mathbf{X}_{m,i} = \tilde{\mathbf{g}}_{m,i} \tilde{\mathbf{v}}_{m,i}^\top$. Based on the definition of our loss function $L'(\cdot)$ in (9), we have that

$$\begin{aligned} \nabla_{x_m} L'(\mathbf{S}_{m,*}) &= \mu^* \mathbf{S}_{m,*} - \frac{2}{\tau_\ell^E} \sum_{i=1}^{\tau_\ell^E} \tilde{\psi}_\nu(r_{m,i} \cdot Q(x_m)) \\ &\stackrel{(a)}{=} \left[\mathbb{E}(r_{m,1} \cdot Q(\mathbf{X}_{m,1})) - \frac{1}{\tau_\ell^E} \sum_{i=1}^{\tau_\ell^E} \tilde{\psi}_\nu(r_{m,i} \cdot Q(\mathbf{X}_{m,i})) \right] \end{aligned}$$

where (a) follows using the same steps as in Lemma 11. Similarly, using the same steps for a single task as in Lemma 11 we have

$$\mathbb{P} \left(\|\nabla L'(\mu^* \mathbf{S}_{m,*})\|_{\text{op}} \geq \frac{2t}{\sqrt{\tau_\ell^E}} \right) \leq 2(k_1 + k_2) \exp \left(-\nu t \sqrt{\tau_\ell^E} + \frac{\nu^2 (4 + S_0^2) C k_1 k_2 \tau_\ell^E}{2} \right)$$

By plugging the values of t and ν in Lemma 11, we finish the proof. \square

Lemma 19. (Concentration of $\hat{\mathbf{S}}_{m,\ell}$) For any low-rank linear model with samples $\mathbf{X}_1 \dots, \mathbf{X}_{\tau_\ell^E}$ drawn from \mathcal{X} according to \mathcal{D} then for the optimal solution to the nuclear norm regularization problem in (2) with $\nu = \sqrt{2 \log(2(k_1 + k_2)/\delta_\ell) / ((4 + S_0^2) \tau_\ell^E k_1 k_2)}$ and

$$\gamma_{m,\ell} = 4 \sqrt{\frac{2(4 + S_0^2) C k_1 k_2 \log(2(k_1 + k_2)/\delta_\ell)}{\tau_{m,\ell}^E}},$$

with probability at least $1 - \delta_\ell$ it holds that:

$$\|\hat{\mathbf{S}}_{m,\ell} - \mu^* \mathbf{S}_{m,*}\|_F^2 \leq \frac{C_1 k_1 k_2 r \log\left(\frac{2(k_1 + k_2)}{\delta_\ell}\right)}{\tau_{m,\ell}^E},$$

for $C_1 = 36(4 + S_0^2)C$, $\|\mathbf{X}\|_F, \|\mathbf{S}_{m,*}\|_F \leq S_0$, some nonzero constant μ^* , and $\mathbb{E}[(\text{SP}(\mathbf{X}))_{ij}^2] \leq C, \forall i, j$. Summing over all phases $\ell \geq 1$ it follows that $\mathbb{P}(\mathcal{F}') \geq 1 - \delta/2$.

Proof. Since the estimator $\hat{\mathbf{S}}_{m,\ell}$ minimizes the regularized loss function defined in Eqn. (6), we have

$$L(\hat{\mathbf{S}}_{m,\ell}) + \gamma_\ell \|\hat{\mathbf{S}}_{m,\ell}\|_{\text{nuc}} \leq L(\mu^* \mathbf{S}_{m,*}) + \gamma_\ell \|\mu^* \mathbf{S}_{m,*}\|_{\text{nuc}}.$$

And due to the fact that $L'(\cdot)$ is a quadratic function, we have the following expression based on multivariate Taylor's expansion:

$$L'(\hat{\mathbf{S}}_{m,\ell}) - L'(\mu^* \mathbf{S}_{m,*}) = \langle \nabla L'(\mu^* \mathbf{S}_{m,*}), \boldsymbol{\Theta} \rangle + 2\|\boldsymbol{\Theta}\|_F^2, \quad \text{where } \boldsymbol{\Theta} = \hat{\mathbf{S}}_{m,\ell} - \mu^* \mathbf{S}_{m,*}.$$

By rearranging the above two results, we can deduce that

$$\begin{aligned} 2\|\boldsymbol{\Theta}\|_F^2 &\leq -\langle \nabla L'(\mu^* \mathbf{S}_{m,*}), \boldsymbol{\Theta} \rangle + \gamma_\ell \|\mu^* \mathbf{S}_{m,*}\|_{\text{nuc}} - \gamma_\ell \|\hat{\mathbf{S}}_{m,*}\|_{\text{nuc}} \\ &\stackrel{(i)}{\leq} \|\nabla L'(\mu^* \mathbf{S}_{m,*})\|_{\text{op}} \|\boldsymbol{\Theta}\|_{\text{nuc}} + \gamma_\ell \|\mu^* \mathbf{S}_{m,*}\|_{\text{nuc}} - \gamma_\ell \|\hat{\mathbf{S}}_{m,\ell}\|_{\text{nuc}}, \end{aligned} \quad (31)$$

where (i) comes from the duality between matrix operator norm and nuclear norm. Next, we represent the saturated SVD of $\mathbf{S}_{m,*}$ as $\mathbf{S}_{m,*} = \mathbf{U} \mathbf{D} \mathbf{V}^\top$ where $\mathbf{U} \in \mathbb{R}^{k_1 \times r}$ and $\mathbf{V} \in \mathbb{R}^{k_2 \times r}$, and here we would work on its full version, i.e.

$$\mathbf{S}_{m,*} = (\mathbf{U}, \mathbf{U}_\perp) \begin{pmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix} (\mathbf{V}, \mathbf{V}_\perp)^\top = (\mathbf{U}, \mathbf{U}_\perp) \mathbf{D}^* (\mathbf{V}, \mathbf{V}_\perp)^\top$$

where we have $\mathbf{U}_\perp \in \mathbb{R}^{k_1 \times (k_1 - r)}$, $\mathbf{D}^* \in \mathbb{R}^{k_1 \times k_2}$ and $\mathbf{V}_\perp \in \mathbb{R}^{k_2 \times (k_2 - r)}$. Furthermore, we define

$$\boldsymbol{\Lambda} = (\mathbf{U}, \mathbf{U}_\perp)^\top \boldsymbol{\Theta} (\mathbf{V}, \mathbf{V}_\perp) = \begin{pmatrix} \mathbf{U}^\top \boldsymbol{\Theta} \mathbf{V} & \mathbf{U}^\top \boldsymbol{\Theta} \mathbf{V}_\perp \\ \mathbf{U}_\perp^\top \boldsymbol{\Theta} \mathbf{V} & \mathbf{U}_\perp^\top \boldsymbol{\Theta} \mathbf{V}_\perp \end{pmatrix} = \boldsymbol{\Lambda}_1 + \boldsymbol{\Lambda}_2$$

where we write

$$\mathbf{\Lambda}_1 = \begin{pmatrix} 0 & 0 \\ 0 & \mathbf{U}_\perp^\top \mathbf{\Theta} \mathbf{V}_\perp \end{pmatrix}, \quad \mathbf{\Lambda}_2 = \begin{pmatrix} \mathbf{U}^\top \mathbf{\Theta} \mathbf{V} & \mathbf{U}^\top \mathbf{\Theta} \mathbf{V}_\perp \\ \mathbf{U}_\perp^\top \mathbf{\Theta} \mathbf{V} & 0 \end{pmatrix}.$$

Afterward, it holds that

$$\begin{aligned} \|\widehat{\mathbf{S}}_\ell\|_{\text{nuc}} &= \|\mu^* \mathbf{S}_{m,*} + \mathbf{\Theta}\|_{\text{nuc}} \stackrel{(a)}{=} \|(\mathbf{U}, \mathbf{U}_\perp) (\mu^* \mathbf{D}^* + \mathbf{\Lambda}) (\mathbf{V}, \mathbf{V}_\perp)^\top\|_{\text{nuc}} \\ &\stackrel{(b)}{=} \|\mu^* \mathbf{D}^* + \mathbf{\Lambda}\|_{\text{nuc}} + \|\mu^* \mathbf{D}^* + \mathbf{\Lambda}_1 + \mathbf{\Lambda}_2\|_{\text{nuc}} \\ &\geq \|\mu^* \mathbf{D}^* + \mathbf{\Lambda}_1\|_{\text{nuc}} - \|\mathbf{\Lambda}_2\|_{\text{nuc}} \\ &= \|\mu^* \mathbf{D}\|_{\text{nuc}} + \|\mathbf{\Lambda}_1\|_{\text{nuc}} - \|\mathbf{\Lambda}_2\|_{\text{nuc}} \\ &= \|\mu^* \mathbf{S}_{m,*}\|_{\text{nuc}} + \|\mathbf{\Lambda}_1\|_{\text{nuc}} - \|\mathbf{\Lambda}_2\|_{\text{nuc}}, \end{aligned} \tag{32}$$

where, (a) follows from the definition of $\mathbf{S}_{m,*}$, and (b) follows the definition of $\mathbf{\Lambda}$. This implies that

$$\|\mu^* \mathbf{S}_{m,*}\|_{\text{nuc}} - \|\widehat{\mathbf{S}}_{m,\ell}\|_{\text{nuc}} \leq \|\mathbf{\Lambda}_2\|_{\text{nuc}} - \|\mathbf{\Lambda}_1\|_{\text{nuc}}.$$

Combining (31) and (32), we have that

$$2\|\mathbf{\Theta}\|_F^2 \leq \left(\|\nabla L(\mu^* \mathbf{B})\|_{\text{op}} + \gamma_\ell\right) \|\mathbf{\Lambda}_2\|_{\text{nuc}} + \left(\|\nabla L(\mu^* \mathbf{B})\|_{\text{op}} - \gamma_\ell\right) \|\mathbf{\Lambda}_1\|_{\text{nuc}}.$$

Then, we refer to the setting in our Lemma 11, and we choose $\gamma_\ell = 4t/\sqrt{M\tau_\ell^E}$ where the value of t is determined in Lemma 11, i.e.

$$\gamma_\ell = 4\sqrt{\frac{2(4 + S_0^2) C k_1 k_2 \log(2(d_1 + d_2)/\delta_\ell)}{M\tau_\ell^E}},$$

we know that $\lambda_{T-1} \geq 2\|\nabla L(\mu^* \mathbf{S}_{m,*})\|_{\text{op}}$ with probability at least $1 - \delta_\ell$ for any $\delta_\ell \in (0, 1)$. Therefore, with a probability at least $1 - \delta_\ell$, we have

$$2\|\mathbf{\Theta}\|_F^2 \leq \frac{3}{2}\gamma_\ell \|\mathbf{\Lambda}_2\|_{\text{nuc}} - \frac{1}{2}\gamma_\ell \|\mathbf{\Lambda}_1\|_{\text{nuc}} \leq \frac{3}{2}\gamma_\ell \|\mathbf{\Lambda}_2\|_{\text{nuc}}.$$

Since we can easily verify that the rank of $\mathbf{\Lambda}_2$ is at most $2r$, and by using Cauchy-Schwarz Inequality we have that

$$2\|\mathbf{\Theta}\|_F^2 \leq \frac{3}{2}\gamma_\ell \sqrt{2r} \|\mathbf{\Lambda}_2\|_F \leq \frac{3}{2}\gamma_\ell \sqrt{2r} \|\mathbf{\Lambda}\|_F = \frac{3}{2}\gamma_\ell \sqrt{2r} \|\mathbf{\Theta}\|_F$$

which implies that

$$\|\mathbf{\Theta}\|_F \leq \frac{3}{4}\sqrt{2r}\gamma_\ell = 6\sqrt{\frac{(4 + S_0^2) C k_1 k_2 r \log\left(\frac{2(k_1 + k_2)}{\delta_\ell}\right)}{\tau_\ell^E}}$$

This implies that $\mathbb{P}(\mathcal{F}'_\ell) \geq 1 - \delta_\ell$. Taking a union bound over all phases $\ell \geq 1$ and recalling $\delta_\ell := \frac{\delta}{2\ell^2}$, we obtain

$$\begin{aligned} \mathbb{P}(\mathcal{F}') &\geq 1 - \sum_{\ell=1}^{\infty} \mathbb{P}((\mathcal{F}'_\ell)^c) \\ &\geq 1 - \sum_{\ell=1}^{\infty} \frac{\delta_\ell}{2} \\ &= 1 - \sum_{\ell=1}^{\infty} \frac{\delta}{4\ell^2} \\ &\geq 1 - \frac{\delta}{2}. \end{aligned}$$

This concludes our proof. \square

A.7 Final Sample Complexity Bound

We first define the arm elimination event similar to Theorem 1. For any $\mathcal{V} \subseteq \underline{\mathcal{W}}$ be the active set and $\underline{\mathbf{w}} \in \mathcal{V}$ define

$$\mathcal{E}_{\underline{\mathbf{w}},\ell}(\mathcal{V}) = \left\{ \left| \left\langle \underline{\mathbf{w}} - \underline{\mathbf{w}}^*, \hat{\boldsymbol{\theta}}_\ell(\mathcal{V}) - \boldsymbol{\theta}^* \right\rangle \right| \leq \epsilon_\ell \right\} \quad (33)$$

where it is implicit that $\hat{\boldsymbol{\theta}}_\ell := \hat{\boldsymbol{\theta}}_\ell(\mathcal{V})$ is the design constructed in the algorithm at stage ℓ with respect to $\underline{\mathcal{W}}_\ell = \mathcal{V}$.

Theorem 2. (Restatement) *With probability at least $1 - \delta$, multi-task **GOBLIN** returns the best arms \mathbf{x}_* , \mathbf{z}_* , and the number of samples used is bounded by*

$$\tilde{O} \left(\frac{M(k_1 + k_2)r}{\Delta^2} + \frac{M\sqrt{k_1 k_2 r}}{S_r} + \frac{\sqrt{d_1 d_2 r}}{S_r} \right).$$

Proof. For the rest of the proof we have that the good events $\mathcal{F}_\ell \cap \mathcal{F}'_\ell \cap \mathcal{E}_{\underline{\mathbf{w}},\ell}(\underline{\mathcal{W}}_\ell)$ holds true for each phase ℓ with probability greater than $(1 - \delta)$. The three events are defined in (25), (29) and (33).

Third Stage: Define $\underline{\mathcal{A}}_{m,\ell} = \{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_\ell : \langle \underline{\mathbf{w}}^* - \underline{\mathbf{w}}, \boldsymbol{\theta}^* \rangle \leq 4\epsilon_{m,\ell}\}$. Note that by assumption $\underline{\mathcal{W}} = \underline{\mathcal{W}}_1 = \underline{\mathcal{A}}_1$. The above lemma implies that with probability at least $1 - \delta$ we have $\bigcap_{\ell=1}^\infty \{\underline{\mathcal{W}}_{m,\ell} \subseteq S_{m,\ell}\}$. This implies that

$$\begin{aligned} \rho^G(\underline{\mathcal{W}}_{m,\ell}) &= \min_{\underline{\mathbf{b}} \in \Delta_{\underline{\mathcal{W}}_m}} \max_{\underline{\mathbf{w}}, \underline{\mathbf{w}}' \in \underline{\mathcal{W}}_{m,\ell}} \|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|_{(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}} \underline{\mathbf{b}}_{\underline{\mathbf{w}} \underline{\mathbf{w}}} \underline{\mathbf{w}}^\top + \boldsymbol{\Lambda})}^{-1} \\ &\leq \min_{\underline{\mathbf{b}} \in \Delta_{\underline{\mathcal{W}}_m}} \max_{\underline{\mathbf{w}}, \underline{\mathbf{w}}' \in S_{m,\ell}} \|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|_{(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}} \underline{\mathbf{b}}_{\underline{\mathbf{w}} \underline{\mathbf{w}}} \underline{\mathbf{w}}^\top + \boldsymbol{\Lambda})}^{-1} \\ &= \rho^G(\underline{\mathcal{A}}_{m,\ell}). \end{aligned}$$

Let the effective dimension be $k = (k_1 + k_2)r$. Define $k_1^\ell = 8k \log(1 + \tau_{m,\ell-1}^G/\lambda)$. For $\ell \geq \lceil \log_2(4\Delta_m^{-1}) \rceil$ we have that $S_{m,\ell} = \{\underline{\mathbf{w}}^*\}$, thus, the sample complexity to identify $\underline{\mathbf{w}}_m^*$ is equal to

$$\begin{aligned} &\sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_m} \left[\tau_{m,\ell}^G \hat{\mathbf{b}}_{m,\ell,\underline{\mathbf{w}}}^G \right] = \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \left(\frac{(k_1^\ell + 1)k_1^\ell}{2} + \tau_{m,\ell}^G \right) \\ &= \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \left(\frac{(k_1^\ell + 1)k_1^\ell}{2} + 2\epsilon_{m,\ell}^{-2} \rho^G(\underline{\mathcal{W}}_{m,\ell}) B_{m,*}^\ell \log(4k_1^\ell \ell^2 |\underline{\mathcal{W}}_m|/\delta) \right) \\ &\stackrel{(a)}{\leq} 2 \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \left(\frac{(k+1)k}{2} \log^2(1 + \tau_{m,\ell-1}^G) + 2\epsilon_{m,\ell}^{-2} \rho^G(\underline{\mathcal{W}}_{m,\ell}) B_{m,*}^\ell \log(4k_1^\ell \ell^2 |\underline{\mathcal{W}}_m|/\delta) \right) \\ &\stackrel{(b)}{\leq} 2 \frac{(k+1)k}{2} \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \left(\log^2(1 + \tau_{m,\ell-1}^G) + 8\epsilon_{m,\ell}^{-2} \rho^G(\underline{\mathcal{W}}_{m,\ell}) B_{m,*}^\ell \log(4k_1^\ell \ell^2 |\underline{\mathcal{W}}_m|/\delta) \right) \\ &\stackrel{(c)}{\leq} (k+1)k \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \left(1 + 16\epsilon_{m,\ell}^{-2} \rho^G(\underline{\mathcal{W}}_{m,\ell}) B_{m,*}^\ell \log^2(4k_1^\ell \ell^2 |\underline{\mathcal{W}}_m|/\delta) \right) \\ &\stackrel{(d)}{\leq} (k+1)k \lceil \log_2(4\Delta_m^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} 32\epsilon_{m,\ell}^{-2} f(\underline{\mathcal{A}}_{m,\ell}) B_{m,*}^\ell \log(4k\ell^2 |\underline{\mathcal{W}}_m|/\delta) \\ &\stackrel{(e)}{\leq} (k+1)k \lceil \log_2(4\Delta_m^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} 32\epsilon_\ell^{-2} f(\underline{\mathcal{A}}_{m,\ell}) (64\lambda S^2 + 64\tau_{m,\ell-1}^G) \log(4k\ell^2 |\underline{\mathcal{W}}|/\delta) \end{aligned}$$

$$\begin{aligned}
&= (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 32\epsilon_\ell^{-2} f(\underline{\mathcal{A}}_{m,\ell}) (64\lambda S^2) \log(4k\ell^2 |\underline{\mathcal{W}}|/\delta) \\
&\quad + (k+1)k \lceil \log_2(4\Delta^{-1}) \rceil + \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} 32\epsilon_\ell^{-2} f(\underline{\mathcal{A}}_{m,\ell}) (64\tau_{m,\ell-1}^G) \log(4k\ell^2 |\underline{\mathcal{W}}|/\delta) \\
&\stackrel{(f)}{\leq} (k+1)k \lceil \log_2(4\Delta_m^{-1}) \rceil + 2048\lambda S^2 \log\left(\frac{4k \log_2^2(8\Delta_m^{-1}) |\underline{\mathcal{W}}_m|}{\delta}\right) \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} 2^{2\ell} f(\underline{\mathcal{A}}_{m,\ell}).
\end{aligned}$$

where, (a) follows as $\log^2(1+\tau_{m,\ell-1}^G/\lambda) \leq \log^2(1+\tau_{m,\ell-1}^G)$, (b) follows by noting that $\log(x \log(1+x)) \leq 2\log(x)$ for any $x > 1$. The (c) follows by subsuming the $\log^2(1+\tau_{m,\ell-1}^G)$ into $2\tau_\ell^G$. The (d) follows as $\log(1+\tau_{m,\ell-1}^G) < \tau_\ell^G$ which enables us to replace the k_1^ℓ inside the log with an additional factor of 2. The (e) follows similarly to (23) by noting that

$$\begin{aligned}
B_{m,*}^\ell &\leq 64(\sqrt{\lambda}S + \sqrt{\lambda_{m,\ell}^\perp} S_{m,\ell}^\perp) \\
&\leq 64\lambda S^2 + \left(\frac{64\tau_{m,\ell-1}^G}{8(d_1+d_2)r \log(1+\frac{\tau_{m,\ell-1}^G}{\lambda})} \right) \cdot \left(\frac{8d_1d_2r}{\tau_\ell^E S_r^2} \log\left(\frac{d_1+d_2}{\delta_\ell}\right) \right) \\
&\stackrel{(a_1)}{\leq} 64\lambda S^2 + 64\tau_{m,\ell-1}^G.
\end{aligned}$$

Finally the (f) follows by subsuming the $\tau_{\ell-1}^G$ with a factor of 2 into the quantity of τ_ℓ^G . Then it follows that

$$\begin{aligned}
\rho_{m,*}^G &= \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}_m}} \max_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_m} \frac{\|\underline{\mathbf{w}} - \underline{\mathbf{w}}^*\|^2 \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_m} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda \right)^{-1}}{(\langle \underline{\mathbf{w}} - \underline{\mathbf{w}}^*, \boldsymbol{\theta}^* \rangle)^2} \\
&= \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}_m}} \max_{\ell \leq \lceil \log_2(4\Delta_m^{-1}) \rceil} \max_{\underline{\mathbf{w}} \in S_{m,\ell}} \frac{\|\underline{\mathbf{w}} - \underline{\mathbf{w}}^*\|^2 \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_m} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda \right)^{-1}}{(\langle \underline{\mathbf{w}} - \underline{\mathbf{w}}^*, \boldsymbol{\theta}^* \rangle)^2} \\
&\geq \frac{1}{\lceil \log_2(4\Delta_m^{-1}) \rceil} \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}_m}} \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \max_{\underline{\mathbf{w}} \in \underline{\mathcal{A}}_{m,\ell}} \frac{\|\underline{\mathbf{w}} - \underline{\mathbf{w}}^*\|^2 \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_m} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda \right)^{-1}}{(\langle \underline{\mathbf{w}} - \underline{\mathbf{w}}^*, \boldsymbol{\theta}^* \rangle)^2} \\
&\geq \frac{1}{16 \lceil \log_2(4\Delta_m^{-1}) \rceil} \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} 2^{2\ell} \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}_m}} \max_{\underline{\mathbf{w}} \in \underline{\mathcal{A}}_{m,\ell}} \frac{\|\underline{\mathbf{w}} - \underline{\mathbf{w}}^*\|^2 \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_m} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda \right)^{-1}}{(\langle \underline{\mathbf{w}} - \underline{\mathbf{w}}^*, \boldsymbol{\theta}^* \rangle)^2} \\
&\geq \frac{1}{64 \lceil \log_2(4\Delta_m^{-1}) \rceil} \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} 2^{2\ell} \inf_{\mathbf{b} \in \Delta_{\underline{\mathcal{W}}_m}} \max_{\underline{\mathbf{w}}, \underline{\mathbf{w}}' \in \underline{\mathcal{A}}_{m,\ell}} \frac{\|\underline{\mathbf{w}} - \underline{\mathbf{w}}'\|^2 \left(\sum_{\underline{\mathbf{w}} \in \underline{\mathcal{W}}_m} \mathbf{b}_{\underline{\mathbf{w}}} \underline{\mathbf{w}} \underline{\mathbf{w}}^\top + \Lambda \right)^{-1}}{(\langle \underline{\mathbf{w}} - \underline{\mathbf{w}}', \boldsymbol{\theta}^* \rangle)^2} \\
&\geq \frac{1}{64 \lceil \log_2(4\Delta_m^{-1}) \rceil} \sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} 2^{2\ell} f(\underline{\mathcal{A}}_{m,\ell}).
\end{aligned}$$

This implies that

$$\sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} 2^{2\ell} f(\underline{\mathcal{A}}_{m,\ell}) \leq \rho_{m,*}^G 64 \lceil \log_2(4\Delta_m^{-1}) \rceil.$$

Plugging this back we get

$$\begin{aligned}
\sum_{\ell=1}^{\lceil \log_2(4\Delta_m^{-1}) \rceil} \sum_{\mathbf{w} \in \mathcal{W}_m} \left[\tau_{m,\ell}^G \widehat{\mathbf{b}}_{\ell,\mathbf{w}} \right] &\leq (k+1)k \lceil \log_2(4\Delta_m^{-1}) \rceil \\
&\quad + 2048\lambda S^2 \log \left(\frac{8k \log_2^2(8\Delta_m^{-1}) |\mathcal{W}_m|}{\delta} \right) 64\rho_{m,*}^G \lceil \log_2(4\Delta_m^{-1}) \rceil \\
&\leq (k+1)k \lceil \log_2(4\Delta_m^{-1}) \rceil \\
&\quad + C_2\lambda S^2 \log \left(\frac{8k \log_2^2(8\Delta_m^{-1}) |\mathcal{W}_m|}{\delta} \right) \rho_{m,*}^G \lceil \log_2(4\Delta_m^{-1}) \rceil
\end{aligned}$$

where, $C_2 > 0$ is a constant. Summing over each task m , the simplified sample complexity for the third stage is given by

$$N_3 \leq O \left(\frac{Mk}{\Delta^2} \log \left(\frac{k \log_2(\Delta^{-1}) |\mathcal{W}_m|}{\delta} \right) \right) = \tilde{O} \left(\frac{M(k_1 + k_2)r}{\Delta^2} \right)$$

where $\Delta = \min_{\mathbf{w} \in \mathcal{W}} (\mathbf{w}_* - \mathbf{w})^\top \boldsymbol{\theta}_* \stackrel{(a1)}{=} \min_{\mathbf{x} \in \mathcal{X} \setminus \{\mathbf{x}_*\}, \mathbf{z} \in \mathcal{Z} \setminus \{\mathbf{z}_*\}} (\mathbf{x}_*^\top \boldsymbol{\Theta}_* \mathbf{z}_* - \mathbf{x}^\top \boldsymbol{\Theta}_* \mathbf{z})$. The (a1) follows by reshaping the arms in \mathcal{W} to recover the arms in \mathcal{X} and \mathcal{Z} .

2nd Stage: Again recall that the E -optimal design in stage 2 of Algorithm 2 satisfies the Assumption 4 as the sample distribution \mathcal{D} has finite second order moments.

For the second stage first observe that by plugging in the definition of $\tilde{\tau}_E^\ell$ we get

$$\begin{aligned}
\|\boldsymbol{\theta}_{m,k+1:p}^*\|_2^2 &= \sum_{i>r \wedge j>r} H_{ij}^2 = \left\| (\widehat{\mathbf{U}}_\ell^\perp)^\top (\mathbf{U}^* \mathbf{S}^* \mathbf{V}^{*\top}) \widehat{\mathbf{V}}_\ell^\perp \right\|_F^2 \\
&\leq \left\| (\widehat{\mathbf{U}}_\ell^\perp)^\top \mathbf{U}^* \right\|_F^2 \|\mathbf{S}^*\|_2^2 \left\| (\widehat{\mathbf{V}}_\ell^\perp)^\top \mathbf{V}^* \right\|_F^2 \leq O \left(\frac{k_1 k_2 r}{\tau_\ell^E} \log \left(\frac{k_1 + k_2}{\delta} \right) \right) \\
&= O \left(\frac{\sqrt{d_1 d_2} r}{S_r} \log \left(\frac{d_1 + d_2}{\delta_\ell} \right) \right),
\end{aligned}$$

which implies $\|\boldsymbol{\theta}_{k+1:p}^*\|_2 = \tilde{O}(\sqrt{k_1 k_2 r}/S_r)$. We also set $\frac{8k_1 k_2 r}{\tau_{m,\ell}^E S_r^2} \log \left(\frac{k_1 + k_2}{\delta_\ell} \right) := S_{m,\ell}^\perp$. Now we bound the sample complexity from the second stage. From the second stage we can show that we have for the arm set \mathcal{W}_m

$$\begin{aligned}
N_2 &= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \sum_{\tilde{\mathbf{w}} \in \tilde{\mathcal{W}}_m} \left[\tilde{\tau}_{m,\ell}^E \widehat{\mathbf{b}}_{m,\ell,\tilde{\mathbf{w}}}^E \right] \\
&= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{(p+1)p}{2} + \tilde{\tau}_{m,\ell}^E \right) \\
&= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{(p+1)p}{2} + \frac{\sqrt{8k_1 k_2 r \log(4\ell^2 |\mathcal{W}|/\delta)}}{S_r} \right) \\
&\leq (p+1)p \lceil \log_2(4\Delta^{-1}) \rceil + 32 \frac{\sqrt{k_1 k_2 r}}{S_r} \log \left(\frac{4 \log_2^2(8\Delta^{-1}) |\mathcal{W}|}{\delta} \right) \lceil \log_2(4\Delta^{-1}) \rceil \\
&= O \left(\frac{\sqrt{k_1 k_2 r}}{S_r} \log \left(\frac{4 \log_2^2(8\Delta^{-1}) |\mathcal{W}|}{\delta} \right) \right) \stackrel{(a)}{=} \tilde{O} \left(\frac{\sqrt{k_1 k_2 r}}{S_r} \right)
\end{aligned}$$

1st Stage: Finally we also use the E -optimal design in first stage of Algorithm 2. Note that this design satisfies the Assumption 4 as the sample distribution \mathcal{D} has finite second-order moments. Now

we bound the sample complexity from the first stage. From the first stage we can show that we have for the arm set $\bar{\mathcal{W}}$

$$\begin{aligned}
N_1 &= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \sum_{m=1}^M \sum_{\mathbf{w} \in \bar{\mathcal{W}}_m} \left[\tau_{m,\ell}^E \hat{\mathbf{b}}_{m,\ell,\mathbf{w}}^E \right] \\
&= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{M(p+1)p}{2} + \sum_m \tau_{m,\ell}^E \right) \\
&= \sum_{\ell=1}^{\lceil \log_2(4\Delta^{-1}) \rceil} \left(\frac{M(p+1)p}{2} + \frac{\sqrt{8d_1d_2r \log(4\ell^2|\mathcal{W}|/\delta)}}{S_r} \right) \\
&\leq M(p+1)p \lceil \log_2(4\Delta^{-1}) \rceil + 32 \frac{\sqrt{d_1d_2r}}{S_r} \log \left(\frac{4\log_2^2(8\Delta^{-1})|\mathcal{W}|}{\delta} \right) \lceil \log_2(4\Delta^{-1}) \rceil \\
&\stackrel{(a)}{=} O \left(\frac{\sqrt{d_1d_2r}}{S_r} \log \left(\frac{4\log_2^2(8\Delta^{-1})|\mathcal{W}|}{\delta} \right) \right) \stackrel{(a)}{\cong} \tilde{O} \left(\frac{\sqrt{d_1d_2r}}{S_r} \right)
\end{aligned}$$

where, (a) follows as $p = d_1d_2$. Combining N_1, N_2 and N_3 gives the claim of the theorem. \square

A.8 Additional Experimental Details

Single Task Unit Ball: This experiment consists of a set of $\{6, 10, 14\}$ left and right arms that are arranged in a unit ball in \mathbb{R}^6 , and $\|\mathbf{x}\| = 1, \|\mathbf{z}\| = 1$ for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$. Hence, we have $d_1 \in \mathbb{R}^6$ and $d_2 \in \mathbb{R}^6$. We choose a random $\Theta_* \in \mathbb{R}^{d_1 \times d_2}$ which has rank $r = 2$. We set $\delta = 0.1$. We compare against RAGE (Fiez et al., 2019) that treats this d_1d_2 bilinear bandit as a linear bandit setting and suffers a sample complexity that scales as $\tilde{O}(d_1d_2/\Delta^2)$. We do a continuous relaxation of the algorithm when implementing it to make this more tractable.

Multi-task Unit Ball: This experiment consists of a set of $\{5, 10, 15, 20, 25, 30\}$ tasks. For each task, we choose left and right arms that are arranged in a unit ball in \mathbb{R}^8 , and $\|\mathbf{x}\| = 1, \|\mathbf{z}\| = 1$ for all $\mathbf{x} \in \mathcal{X}$ and $\mathbf{z} \in \mathcal{Z}$. Hence, we have $d_1 \in \mathbb{R}^8$ and $d_2 \in \mathbb{R}^8$. We choose $k_1 = k_2 = 4$, and feature extractors $\mathbf{B}_1 \in \mathbb{R}^{d_1 \times k_1}, \mathbf{B}_2 \in \mathbb{R}^{d_2 \times k_2}$ shared across tasks. We choose a random matrix $\mathbf{S}_{m,*} \in \mathbb{R}^{k_1 \times k_2}$ for each task m such that $\mathbf{S}_{m,*}$ has rank $r = 2$. We set $\delta = 0.1$. We compare against DouExpDes (Du et al., 2023) that treats this setting as M k_1k_2 bilinear bandits (after learning the feature extractors) and suffers a sample complexity that scales as $\tilde{O}(Mk_1k_2/\Delta^2)$ (see Remark 2). Again we do a continuous relaxation of the algorithm when implementing it to make this more tractable.

B Table of Notations

Notations	Definition
\mathcal{X}	Left arm set
\mathcal{Z}	Right arm set
M	Number of tasks
ℓ	Phase number
$\Theta_{m,*}$	Hidden parameter matrix for
\mathbf{b}_ℓ^E	E-optimal design at the ℓ -th phase
$\mathbf{b}_{m,\ell}^G$	G-optimal design at the ℓ -th phase for the m -th task
$S_{m,\ell}^\perp$	$\frac{8d_1d_2r}{\tau_\ell^E S_r^2} \log\left(\frac{d_1+d_2}{\delta_\ell}\right)$
$\lambda_{m,\ell}^\perp$	$\tau_{m,\ell-1}^G / 8(d_1 + d_2)r \log(1 + \frac{\tau_{m,\ell-1}^G}{\lambda})$
$B_{m,*}^\ell$	$(8\sqrt{\lambda}S + \sqrt{\lambda_{m,\ell}^\perp S_{m,\ell}^\perp})$
\mathbf{B}_1	Left feature extractor
\mathbf{B}_2	Right feature extractor
S_r	r -th largest singular value of Θ_*
$\Delta(\mathbf{x}, \mathbf{z})$	$\mathbf{x}_*^\top \Theta_* \mathbf{z}_* - \mathbf{x}^\top \Theta_* \mathbf{z}$
Δ	$\min_{\mathbf{x} \neq \mathbf{x}_*, \mathbf{z} \neq \mathbf{z}_*} \Delta(\mathbf{x}, \mathbf{z})$
$\mathcal{Y}(\mathcal{W})$	$\{\mathbf{w} - \mathbf{w}' : \forall \mathbf{w}, \mathbf{w}' \in \mathcal{W}, \mathbf{w} \neq \mathbf{w}'\}$
$\mathcal{Y}^*(\mathcal{W})$	$\{\mathbf{w}_* - \mathbf{w} : \forall \mathbf{w} \in \mathcal{W} \setminus \mathbf{w}_*\}$
δ	confidence level

Table 1: Table of Notations