
Multi-task Representation Learning with Stochastic Linear Bandits

Leonardo Cella

CSML, Italian Institute of Technology

Karim Lounici

CMAP, Ecole Polytechnique

Grégoire Pacreau

CMAP, Ecole Polytechnique

Massimiliano Pontil

CSML, Italian Institute of Technology
& Dept of Computer Science, UCL

Abstract

We study the problem of transfer-learning in the setting of stochastic linear contextual bandit tasks. We consider that a low dimensional linear representation is shared across the tasks, and study the benefit of learning the tasks jointly. Following recent results to design Lasso stochastic bandit policies, we propose an efficient greedy policy based on trace norm regularization. It implicitly learns a low dimensional representation by encouraging the matrix formed by the task regression vectors to be of low rank. Unlike previous work in the literature, our policy does not need to know the rank of the underlying matrix, nor does it require the covariance of the arms distribution to be invertible. We derive an upper bound on the multi-task regret of our policy, which is, up to logarithmic factors, of order $T\sqrt{rN} + \sqrt{rNTd}$, where T is the number of tasks, r the rank, d the number of variables and N the number of rounds per task. We show the benefit of our strategy over an independent task learning baseline, which has a worse regret of order $T\sqrt{dN}$. We also argue that our policy is minimax optimal and, when $T \geq d$, has a multi-task regret which is comparable to the regret of an oracle policy which knows the true underlying representation.

1 INTRODUCTION

Contextual bandits (Abbasi-Yadkori et al., 2011; Li et al., 2010; Auer, 2002) are a prominent learning framework to study sequential decision problems with partial feedback.

They find applications in numerous fields, ranging from recommender systems (Li et al., 2010), to finance (Shen et al., 2015) and to adaptive routing (Awerbuch and Kleinberg, 2008), among others. This methodology was originally motivated by applications in clinical trials (Woodroffe, 1979), whereby a doctor has to decide which among available treatments is best suited for a patient, through a sequence of trials. A fundamental aspect in bandit problems is to control the trade-off between *exploration* and *exploitation*, namely, the balance between the need of acquiring more information and the temptation to act optimally according to the already available knowledge.

In this paper we study multi-task learning with stochastic linear contextual bandit tasks (Lu et al., 2021; Li et al., 2017; Filippi et al., 2010). Within this setting, each task is associated with a regression vector and proceeds sequentially. At each trial an agent observes a set of different alternatives (*arms*) linked to a feature (*context*) vector. The agent then selects one context vector and subsequently observes a stochastic reward generated by a noisy linear regression associated to the chosen vector. The goal is to design an algorithm (policy) that learns, interacting with the tasks, how to select contexts that are most aligned with the underlying task (regression vector), hence maximizing the cumulative rewards across all the tasks.

A central idea of multi-task learning is to leverage similarities between tasks in order to facilitate learning. In this paper, we consider that the tasks share a low dimensional representation, that is, the task regression vectors span a low dimensional subspace. The benefit of learning such representation has been widely investigated in both the standard supervised learning and in the reinforcement learning settings (see Lounici et al., 2011; Koltchinskii et al., 2011; Negahban and Wainwright, 2011; Calandriello et al., 2014, and references therein). In the bandit setting, this problem presents additional difficulties, since contexts vectors are no longer independently distributed. Indeed, they are collected sequentially depending on past observations and on the adopted bandit policy. This entails two main challenges.

On the one hand, we would like the collected contexts to span the whole feature space, as this would facilitate the estimation of the unknown regression vectors. On the other hand, collecting contexts which are misaligned with respect to the regression vector results in poor performance for the bandit strategy.

Contributions. In this work we present an efficient policy based on the trace norm regularization estimator (see Koltchinskii et al., 2011; Negahban and Wainwright, 2011; Maurer and Pontil, 2013, and references therein) which leverages the tasks common representation to improve learning. We provide oracle inequalities for this policy under the restricted strong convexity condition with correct and explicit dependencies (Lemma 1 and Proposition 1) that are based on a novel martingale concentration argument. Next, we provide an upper bound for the proposed policy (Theorem 1) which is valid without boundedness assumption on the arms and is minimax optimal (up to a logarithmic factor). We noticed that when the number of tasks is larger than the ambient dimension our policy is comparable to the oracle policy which knows the true underlying representation a-priori. A key novelty of the proposed policy over the state-of-the-art (Yang et al., 2020) is that it does not need to know the rank of the representation matrix, it does not require the arms covariance distribution to be invertible and its regret bound is non-trivial already when the time horizon is $O(d)$ in the worse case and potentially $O(\log d)$ in favorable scenarios.

Notation. For a real vector $\mathbf{x} \in \mathbb{R}^d$, we use $\|\mathbf{x}\|$ to denote its Euclidean norm. Given a pair of symmetric matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{d \times d}$, the expression $\mathbf{A} \succeq \mathbf{B}$ means that $\mathbf{A} - \mathbf{B}$ is positive semi-definite. We respectively use $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ to refer to the minimum and maximum eigenvalues of a square symmetric matrix \mathbf{A} . Similarly, we respectively use $\sigma_{\min}(\mathbf{A})$ and $\sigma_{\max}(\mathbf{A})$ for the smallest and the largest singular values of a generic matrix \mathbf{A} . Given a positive definite matrix $\mathbf{A} \succ \mathbf{0}$, we indicate with $\|\mathbf{x}\|_{\mathbf{A}} = \sqrt{\mathbf{x}^T \mathbf{A} \mathbf{x}}$ the corresponding weighted Euclidean norm. For any matrix $\mathbf{A} \in \mathbb{R}^{d \times T}$, we let $\|\mathbf{A}\|_* = \sum_{i=1}^d \sigma_i(\mathbf{A})$ be the trace norm (sum of its singular values), $\|\mathbf{A}\|_F$ the Frobenius norm and $\|\mathbf{A}\|_{\text{op}} = \sigma_{\max}(\mathbf{A})$ the operator norm (largest singular value). We denote with $[\mathbf{A}]^i$ its i -th row and with $[\mathbf{A}]_j$ its j -th column. Finally, we use $\text{diag}(\lambda_1, \dots, \lambda_d)$ for the $d \times d$ diagonal matrix with values $\lambda_1, \dots, \lambda_d$ on the diagonal. Given a random event Υ , we denote with Υ^c its complement. Finally, given a scalar $\epsilon \in \mathbb{R}$ we denote with ϵ_d the d -dimensional vector having value ϵ in each component. Additional notation is introduced on the way; a table summarizing the notation used throughout the paper is reported in the appendix.

2 RELATED WORKS

In the last two decades many efforts have been devoted to designing contextual bandit policies (Abbasi-Yadkori et al., 2011, 2012; Ariu et al., 2020; Auer, 2002; Bastani and Bayati, 2020; Chu et al., 2011; Kim and Paik, 2019; Li et al., 2017; Oh et al., 2021; Wang et al., 2018; Kuzborskij et al., 2019; Foster and Rakhlin, 2020). When restricting to the high-dimensional setting an appealing approach is based on sparse linear models, i.e. the number s of non-zero components of the regression vector is assumed to be much smaller than the input dimension d . Many different strategies have been investigated (Kim and Paik, 2019; Ariu et al., 2020; Bastani and Bayati, 2020). Among the proposed approaches, one of the most recent and valuable works is (Oh et al., 2021), where they design a greedy policy based on the Lasso estimator. Interestingly, their approach does not require knowledge of s and it does not perform random pulls in order to have i.i.d. data. Inspired by this work, we propose a greedy algorithm which does not need to know the rank index r (dimensionality of the underlying representation). A technical challenge that we are facing is how to obtain an accurate estimator relying on non i.i.d. samples, while considering a more complex matrix estimator. We observe that their regret bound argument is not accurate as there might be a hidden dependency in the number of features (the same inaccuracy can be found in (Cella and Pontil, 2021; Kim and Paik, 2019; Calandriello et al., 2014)).

Multi-task and meta-learning frameworks have been studied primarily in the supervised-learning setting (Tripuraneni et al., 2021; Denevi et al., 2018, 2019; Lounici et al., 2011; Argyriou et al., 2008; Baxter, 2000). Specifically, the impact of representation learning with trace-norm regularization has been widely investigated when considering i.i.d. data (Koltchinskii et al., 2011; Negahban and Wainwright, 2011; Maurer and Pontil, 2013). More recently different authors have investigated the combination of multi-task and meta-learning with interactive learning settings (e.g. bandits and reinforcement learning) (Hu et al., 2021; Basu et al., 2021; Cella and Pontil, 2021; Kveton et al., 2021; Simchowit et al., 2021; Calandriello et al., 2014; D’Eramo et al., 2019; Cella et al., 2020; Yang et al., 2020). Among the latter category, the most relevant works are (Yang et al., 2020; Hu et al., 2021; Cella and Pontil, 2021). Cella and Pontil (2021) considered both the multi-task and the meta-learning frameworks but assuming the task vector parameters to be *jointly sparse*, which is more restrictive than the low rank assumption considered here. Differently, (Yang et al., 2020) considers a similar multi-task problem both in the finite and infinite action settings, whereas (Hu et al., 2021) considers only the latter setting. Note that these two references impose stronger conditions and assumptions. Indeed, the policies considered in those papers require knowledge of the low-rank parameter r and are not sample efficient. In this work, we develop a different and more general analy-

sis for the finite action setting. Notably, our algorithm is computationally efficient and does not require knowledge of low-rank parameter r . Our idea is similar to the last approach investigated in (Calandriello et al., 2014), but their work considers the simpler case where samples are independently distributed. At last we mention that (Kveton et al., 2017) studied low-rank stochastic bandits. However they consider a completely different protocol, regret and arms definitions. Thus it is not possible to draw a quantitative comparison to this work.

3 PROBLEM SETTING

In this section we recall the stochastic linear contextual bandit problem and then introduce the multi-task learning framework.

3.1 Stochastic Linear Contextual Bandits

The contextual linear bandit problem consists of a sequence of N interactions between a learning policy π and the environment. At each round $n \in [N]$, the agent is given a set $\mathcal{D}_n \subseteq \mathbb{R}^d$ formed of K arms (context vectors) from which it has to select one arm $\mathbf{x}_n \in \mathcal{D}_n$ among the available ones. Then, the agent observes the corresponding reward Y_n via a noisy linear regression,

$$Y_n = \mathbf{x}_n^\top \mathbf{w} + \eta_n, \quad (1)$$

where η_n is a random variable whose statistical distribution will be specified later. By knowing the true regression vector \mathbf{w} , at each round $n \in [N]$ the optimal policy π^* selects $\mathbf{x}_n^* = \arg \max_{\mathbf{x} \in \mathcal{D}_n} \mathbf{x}^\top \mathbf{w}$, maximizing the instantaneous expected reward. The learning objective is then to maximize the cumulative reward, or equivalently, to minimize the *pseudo-regret*

$$R(N, \mathbf{w}) = \sum_{n=1}^N r_n = \sum_{n=1}^N (\mathbf{x}_n^* - \mathbf{x}_n)^\top \mathbf{w}.$$

We make the following additional assumptions on the arms, the regression vector and the noise variables.

Assumption 1 (Contexts Distribution). *At each round $n \in [N]$, the decision set $\mathcal{D}_n \subset \mathbb{R}^d$ consists of K d -dimensional vectors \mathbf{x}_k admitting representation $\mathbf{x}_k = \Sigma_k^{1/2} \mathbf{z}_k$ where Σ_k is the covariance operator of \mathbf{x}_k and \mathbf{z}_k is a d -dimensional vector with independent sub-Gaussian entries with zero mean and variance 1. We assume the tuples $\mathcal{D}_1, \dots, \mathcal{D}_N$ to be drawn i.i.d. from a fixed unknown zero mean sub-Gaussian joint distribution p on \mathbb{R}^{Kd} .*

Let $C_{\mathbf{z}}$ be a positive constant satisfying

$$\max_{1 \leq k \leq K} \|\mathbf{z}_k\|_{\psi_2} < C_{\mathbf{z}}.$$

We denote by $C_{\mathbf{x}}$ the sub-Gaussian Orlicz norm of the K marginal distributions of p corresponding to the K arms,

that is $C_{\mathbf{x}} := \max_{1 \leq k \leq K} \|\mathbf{x}_k\|_{\psi_2}$. An obvious computation gives

$$C_{\mathbf{x}} \leq C_{\mathbf{z}} \max_{1 \leq k \leq K} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} < \infty.$$

The above assumption is quite standard when considering high-dimensional linear bandits (Bastani and Bayati, 2020; Hao et al., 2020; Kim and Paik, 2019; Rusmevichientong and Tsitsiklis, 2010). Note that vectors associated to different arms are allowed to be correlated between each other. Moreover, Assumption 1 implies that each arm k admits zero mean, square-integrable marginal distribution with $d \times d$ covariance Σ_k .

Finally, similarly to (Bastani et al., 2021; Oh et al., 2021), we introduce the following assumption on the arms distribution. This mild condition is necessary for our analysis in order to control the fulfillment of a specific regularity property (the RSC condition, see Definition 1 below) by the empirical covariance matrix. Specifically, it will allow the designed policy to avoid interleaving its arm selection strategy with random plays.

Assumption 2 (Arms distribution). *There exists a constant $\nu < \infty$ such that $p(-\bar{\mathbf{x}})/p(\bar{\mathbf{x}}) \leq \nu \forall \bar{\mathbf{x}} \in \mathbb{R}^{Kd}$. Moreover, there exists a constant $\omega_{\mathcal{X}} < \infty$ such that, for any permutation (a_1, \dots, a_K) of $[K]$, any integer $i \in \{2, \dots, K-1\}$ and any fixed vector $\mathbf{w} \in \mathbb{R}^d$, it holds that*

$$\begin{aligned} \mathbb{E} [\mathbf{x}_{a_i} \mathbf{x}_{a_i}^\top \mathbb{I} \{ \mathbf{x}_{a_1}^\top \mathbf{w} < \dots < \mathbf{x}_{a_K}^\top \mathbf{w} \}] &\preceq \\ &\preceq \omega_{\mathcal{X}} \mathbb{E} [(\mathbf{x}_{a_1} \mathbf{x}_{a_1}^\top + \mathbf{x}_{a_K} \mathbf{x}_{a_K}^\top) \mathbb{I} \{ \mathbf{x}_{a_1}^\top \mathbf{w} < \dots < \mathbf{x}_{a_K}^\top \mathbf{w} \}]. \end{aligned}$$

Parameter ν characterizes the skewness of the arms distribution p ; for symmetrical distributions $\nu = 1$. Notice that Assumption 2 is satisfied for a large class of distributions both discrete and continuous (e.g. Gaussian, multi-dimensional uniform and Rademacher). The value $\omega_{\mathcal{X}}$ captures dependencies between arms: the more positively correlated they are, the smaller $\omega_{\mathcal{X}}$ will be. The extreme scenario is given by perfectly correlated arms, in which case we have $\omega_{\mathcal{X}}$ independent of any problem parameters. Finally, when arms are generated i.i.d. from a multivariate Gaussian or a multivariate uniform distribution over the sphere we have $\omega_{\mathcal{X}} = O(1)$.

3.2 Low-Rank Linear Contextual Bandits

In this paper, we address the problem of simultaneously solving T linear contextual bandit tasks. Each task $t \in [T]$ lasts for N rounds and is associated with a regression vector $\mathbf{w}_t \in \mathbb{R}^d$. We denote with

$$\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_T]$$

the $d \times T$ matrix, whose columns are formed by the T regression vectors, which are unknown to the learner. The bandit tasks are explored in parallel. At each round $n \in [N]$

each task simultaneously receives a decision set $\mathcal{D}_{t,n}$ from the environment, upon which the learner (policy π) picks an arm $\mathbf{x}_{t,n}$, and the associated feedback $y_{t,n}$ is observed via the linear regression

$$y_{t,n} = \mathbf{x}_{t,n}^\top \mathbf{w}_t + \eta_{t,n} \quad (2)$$

where $\eta_{t,n}$ is a noise random variable which we specify below.

We make the following assumption on the task regression vectors.

Assumption 3 (Low-Rank Assumption). *The matrix $\mathbf{W} \in \mathbb{R}^{d \times T}$ has rank $\rho(\mathbf{W}) = r$, with $r \ll \min(d, T)$.*

The above assumption implies that there exists a low dimensional representation $\mathbf{B} \in \mathbb{R}^{d \times r}$ with orthonormal columns and a matrix $\mathbf{C} \in \mathbb{R}^{r \times T}$ such that $\mathbf{W} = \mathbf{B}\mathbf{C}$. This is in line with the standard high-dimensional setting where many features are redundant.

Our principal objective is to minimize the multi-task pseudo-regret,

$$\begin{aligned} \bar{R}(T, N) &= \sum_{t=1}^T R(N, \mathbf{w}_t) \\ &= \sum_{t=1}^T \sum_{n=1}^N (\mathbf{x}_{t,n}^* - \mathbf{x}_{t,n})^\top \mathbf{w}_t, \end{aligned} \quad (3)$$

where $\mathbf{x}_{t,n}^* = \arg \max_{\mathbf{x} \in \mathcal{D}_{t,n}} \mathbf{x}^\top \mathbf{w}_t$.

4 LOW-RANK MATRIX ESTIMATION

In the standard supervised learning setting a natural estimator suited to Assumption 3 is given by the trace (nuclear) norm regularized estimator (see [Argyriou et al., 2008](#); [Bühlmann and Van De Geer, 2011](#); [Lounici et al., 2011](#); [Negahban and Wainwright, 2011](#), and references therein). In particular a large body of works have shown that bounds on the estimation error is controlled by the rank of the underlying regression matrix. Here we adapt this methodology to the bandit setting. At each round $n \in [N]$, we estimate matrix \mathbf{W} via $\widehat{\mathbf{W}}_{n+1}$ as the solution of the following trace norm regularization problem

$$\arg \min_{\mathbf{A} \in \mathbb{R}^{d \times T}} \frac{1}{n} \sum_{t=1}^T \|\mathbf{y}_{t,n} - \mathbf{X}_{t,n}[\mathbf{A}]_t\|_2^2 + \lambda_n \|\mathbf{A}\|_* \quad (4)$$

where the design matrix $\mathbf{X}_{t,n} \in \mathbb{R}^{n \times d}$ contains the context vectors $\mathbf{x}_{t,i} \in \mathbb{R}^d$, $i \in [n]$ as its rows, the vector $\mathbf{y}_{t,n} \in \mathbb{R}^n$ is formed by the rewards for task t after n interactions, sampled from (2), and $\|\mathbf{A}\|_*$ is the trace norm of the matrix \mathbf{A} , that is the sum of its singular values. If compared to the Lasso estimator, the objective function (4) encourages low-rank matrices instead of sparse vectors. Before presenting

the technical results we need to introduce the following additional notation.

Covariance matrices. We indicate the theoretical averaged $d \times d$ covariance matrix as

$$\Sigma = \frac{1}{K} \mathbb{E} \left[\sum_{k=1}^K \mathbf{x}_k \mathbf{x}_k^\top \right] = \frac{1}{K} \sum_{k=1}^K \Sigma_k, \quad (5)$$

where the expectation is over the decision set sampling distribution $p_{\mathcal{X}}$ which is assumed to be shared between the tasks.

For every $t \in [T]$, we denote the empirical covariance matrix for task t as

$$\widehat{\Sigma}_{t,n} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_{t,i} \mathbf{x}_{t,i}^\top. \quad (6)$$

Moreover, we use the notation $\bar{\Sigma}$ and $\widehat{\Sigma}_n \in \mathbb{R}^{dT \times dT}$ for the theoretical and the empirical multi-task matrices respectively. They are both block diagonal and composed by the corresponding T single task $d \times d$ matrices on the diagonal, that is

$$\bar{\Sigma} = \text{diag}(\Sigma, \dots, \Sigma)$$

and

$$\widehat{\Sigma}_n = \text{diag}(\widehat{\Sigma}_{1,n}, \dots, \widehat{\Sigma}_{T,n}).$$

We introduce now the restricted strong convexity (RSC) condition on covariance $\bar{\Sigma}$. To this end, we denote by $\text{Vec}(\Delta)$, the vector in \mathbb{R}^{dT} obtained by stacking together the columns of $\Delta \in \mathbb{R}^{d \times T}$. Let

$$\mathbf{W} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$$

be the singular value decomposition of matrix \mathbf{W} of rank r where $\mathbf{U} \in \mathbb{R}^{d \times r}$, $\mathbf{V} \in \mathbb{R}^{T \times r}$ are the matrices formed by the left and right singular vectors, respectively and \mathbf{D} is the $r \times r$ diagonal matrix of singular values.

Definition 1 (RSC Condition). *We say that the restricted strong convexity (RSC) condition is met for the theoretical multi-task matrix $\bar{\Sigma} \in \mathbb{R}^{dT \times dT}$, with positive constant $\kappa(\bar{\Sigma})$ if*

$$\min \left\{ \frac{\|\text{Vec}(\Delta)\|_{\bar{\Sigma}}^2}{2 \|\text{Vec}(\Delta)\|_2^2} : \Delta \in \mathcal{C}(r) \right\} \geq \kappa(\bar{\Sigma}), \quad (7)$$

where

$$\mathcal{C}(r) = \left\{ \Delta \in \mathbb{R}^{d \times T} : \|\Pi(\Delta)\|_* \leq 3 \|\Delta - \Pi(\Delta)\|_* \right\} \quad (8)$$

and $\Pi(\Delta)$ is the projection onto set

$$\left\{ \Delta \in \mathbb{R}^{d \times T} : \text{Col}(\Delta) \perp \mathbf{U}, \text{Row}(\Delta) \perp \mathbf{V} \right\}.$$

The RSC condition allows us to control the error $\Delta_n = \widehat{\mathbf{W}}_n - \mathbf{W}$ as it guarantees that the considered empirical loss is strictly convex on a restricted subset of approximately low rank matrices defined by the cone $\mathcal{C}(r)$.

Remark 1 (Value of $\kappa(\Sigma)$). As discussed in (Calandriello et al., 2014) Definition 1 can be compared with the restricted eigenvalue and the compatibility conditions that have been investigated for the group Lasso (Bühlmann and Van De Geer, 2011; Lounici et al., 2011) and Lasso (Oh et al., 2021; Kim and Paik, 2019; Calandriello et al., 2014) estimators. It is standard in high-dimensional statistics or in compressed sensing to assume the existence of an absolute constant $\kappa > 0$ such that $\kappa(\widehat{\Sigma}) > \kappa > 0$. Such conditions are satisfied w.h.p. for instance by context vectors with i.i.d. zero mean, variance 1, sub-Gaussian entries.

4.1 Oracle Inequality with non i.i.d. Data

We can now state our first result which controls the Forbeniou-norm estimation error

$$\left\| \widehat{\mathbf{W}}_{n+1} - \mathbf{W} \right\|_{\text{F}}$$

assuming the RSC condition to hold for the empirical multi-task matrix $\widehat{\Sigma}_n$. In Proposition 2 we will show that such condition is satisfied with high probability under our Assumptions 1, 2 and the RSC condition on $\widehat{\Sigma}$.

We denote by $\{\mathbf{e}_1, \dots, \mathbf{e}_T\}$ the standard canonical basis of \mathbb{R}^T . We have then to upper bound the operator norm of the following matrix

$$\mathbf{D}_n = \sum_{t=1}^T \sum_{i=1}^n \eta_{t,i} \mathbf{x}_{t,i} \mathbf{e}_t^\top$$

in order to set the regularization parameter at round n so that the estimation bound in the following lemma holds with high probability.

Lemma 1. Let $\{\mathbf{x}_{t,i} : t \in [T], i \in [N]\}$ be the sequence generated by Algorithm (1) Suppose the RSC condition holds for the empirical multi-task matrix $\widehat{\Sigma}_n$ with constant $\kappa(\widehat{\Sigma}_n)$. For $\delta \in (0, 1)$, define the regularization parameter $\lambda_n = \lambda_n(\delta)$ such that with probability at least $1 - \delta$

$$\frac{1}{n} \|\mathbf{D}_n\|_{\text{op}} \leq \lambda_n. \quad (9)$$

Then with probability at least $1 - \delta$ the trace-norm regularized estimate $\widehat{\mathbf{W}}_n$ defined in (4) satisfies

$$\left\| \widehat{\mathbf{W}}_{n+1} - \mathbf{W} \right\|_{\text{F}} \leq \frac{32\lambda_n \sqrt{r}}{\kappa(\widehat{\Sigma}_n)}. \quad (10)$$

The proof, which is presented in the appendix, follows along the lines of (Negahban and Wainwright, 2011, Theorem 1).

Remark 2. We note that a similar result can be also be found in Koltchinskii et al. (2011). In addition, as pointed out in these references, it is possible to extend this result to the case of approximately low rank matrices \mathbf{W} by introducing an additional misspecification error in (10).

4.2 Controlling the Noise Term

The result in Lemma 1 requires us to set the regularization parameter at round n as in (9). To this end we exploit deviation bounds for martingales. Let us define the filtration $(\overline{\mathcal{F}}_n)_{n \geq 0}$ on a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ as follows: $\overline{\mathcal{F}}_0$ is the trivial σ -field $\{\emptyset, \Omega\}$, and for any $n \geq 1$,

$$\overline{\mathcal{F}}_n = \sigma(\mathbf{X}_{1,i}, \eta_{1,i}, \dots, \mathbf{X}_{T,i}, \eta_{T,i}, i \in [n]).$$

We consider the standard noise assumption (Abbasi-Yadkori et al., 2011; Bastani and Bayati, 2020; Cella et al., 2020).

Assumption 4 (subGaussian noise). The noise variables $(\eta_{t,n})_{t,n}$ are a sequence of sub-Gaussian random variables adapted to the filtration $\{\overline{\mathcal{F}}_n\}_{n \geq 0}$ and such that for any $1 \leq t \leq T$ and $n \geq 1$,

$$\mathbb{E}[\eta_{t,n} | \overline{\mathcal{F}}_{n-1}] = 0, \quad \text{and} \quad \mathbb{E}[\eta_{t,n}^2 | \overline{\mathcal{F}}_{n-1}] \leq \sigma^2,$$

and $\eta_{1,n}, \dots, \eta_{T,n}$ are mutually independent conditionally on $\overline{\mathcal{F}}_{n-1}$. We denote by c_η the sub-Gaussian norm of the $\eta_{t,n}$'s, that is, $\max_{t,n} \{\|\eta_{t,n}\|_{\psi_2}\} \leq c_\eta$.

Proposition 1 (SubGaussian noise). Let Assumptions 1 and 4 be satisfied. Then, with probability at least $1 - \delta$

$$\frac{1}{n} \|\mathbf{D}_n\|_{\text{op}} \leq \lambda_n$$

where

$$\lambda_n = \max_{k \in [K]} \|\Sigma_k\|_{\text{op}}^{1/2} \left[\sigma \sqrt{\frac{(d+T) \log \frac{2N(d+T)}{\delta}}{n}} \sqrt{\frac{C \sqrt{(T+d + \log \frac{4N}{\delta}) \log^3 \frac{8N(d+T)}{\delta}}}{n}} \right] \quad (11)$$

where $C > 0$ can depend only on $c_\eta, C_{\mathbf{z}}$.

Proof Sketch. The complete proof is given in Appendix B. The proof relies on Freedman's inequality for martingale (Tropp, 2011, Corollary 1.3) combined with a stopping time argument to handle unbounded arms. \square

5 MULTI-TASK BANDITS VIA TRACE NORM REGULARIZATION

In this section we present our proposed algorithm for multi-task learning with linear stochastic bandits. The algorithm

Algorithm 1 Trace-Norm Bandit

Require: Confidence parameter δ , noise variance σ^2

- 1: At round $n = 1$ arms are picked randomly
- 2: Observe $Y_{1,1}, \dots, Y_{T,1}$
- 3: **for** $n \in 2, \dots, N$ **do**
- 4: update $\widehat{\mathbf{W}}_n$ and λ_n according to (4) and (11)
- 5: **for** $t \in 1, \dots, T$ **do**
- 6: observe $\mathcal{D}_{t,n}$
- 7: pick $\mathbf{x}_{t,n} \in \arg \max_{\mathbf{x} \in \mathcal{D}_{t,n}} \mathbf{x}^\top [\widehat{\mathbf{W}}_n]_t$
- 8: observe reward $Y_{t,n}$
- 9: **end for**
- 10: **end for**

relies on the trace-norm regularized estimator in (4) and is displayed in Algorithm 1. It does not require to know any input parameter besides an upper bound on the noise variance σ^2 and the confidence value δ , which are necessary to specify the regularization parameter¹ λ_n . Indeed this is the only parameter required by our policy and it should be chosen large enough (in theory) as prescribed by Proposition 1; see also equation (11).

Algorithm 1 updates the task parameters via a trace-norm regularized estimation using all the already observed data. Similarly to (Oh et al., 2021), our policy is completely greedy and does not need any forced-sampling strategy. This phenomenon was recently highlighted by (Bastani et al., 2021) which proved that greedy policies do not need an explicit exploration and can even be rate-optimal if there is sufficient randomness in the observed contexts (relying on i.i.d. arms assumption). This is unlike the UCB-like policies (like OFUL) or forced-exploration policies Bastani and Bayati (2020). Our trace-norm bandit policy (Alg. 1) belongs to the former class of “exploration-free” greedy policies. More precisely, our policy performs a natural initial exploration thanks to the arms randomness condition (Assumptions 1 and 2) until the empirical covariance matrix satisfies the (RSC) condition. In the regret analysis below this can be seen in the definition of the first N_0 rounds in (12).

5.1 Multi-Task Learning

In this subsection we present our main result which is a high-probability upper bound on the multi-task regret incurred by the trace-norm bandit policy.

For any $\delta \in (0, 1)$, let $N_0(\delta)$ be the smallest integer such that

$$N \geq C \max_{k \in [K]} \|\Sigma_k\|_{\text{op}}^2 \left((r \log d) \log \frac{4TN}{\delta} \right)^2, \quad (12)$$

¹Using a regularization is quite common in the bandit literature (e.g. Abbasi-Yadkori et al. (2011); Bastani and Bayati (2020).

where $C > 0$ is some large enough numerical constant that may depend only on the distribution of the arms, in particular $\nu, \omega_{\mathcal{X}}$ and $C_{\mathbf{z}}$.

Note that (12) is a sufficient condition on the minimum number of rounds for our regret bound to be valid. Interestingly it depends on the ambient dimension d only logarithmically. This means that for our greedy policy, we can guarantee that the duration of the implicit exploration phase is at most r^2 (up to logarithmic factors). In situation where the \mathbf{W}^* matrix is low-rank with $r \ll \sqrt{d}$, this represents another benefit of our trace norm policy procedure over (Yang et al., 2020) which requires a much longer exploration phase of at least d^2 rounds for their regret bound to be valid.

We now state the main result.

Theorem 1. *Let Assumptions 1, 2, 3, 4 be satisfied. Assume that $\max_{t \in [T]} \|\mathbf{w}_t\|_2 \leq L$, for some constant $L > 0$ and that $N \geq N_0(\delta)$ for some $\delta \in (0, 1)$. Then, with probability at least $1 - \delta$, the multi-task regret of Algorithm 1 is upper bounded (up to logarithmic factors) by*

$$\mathcal{O} \left((T\sqrt{rN} + \sqrt{rdTN}) \right).$$

Our proof is inspired by the one proposed in (Oh et al., 2021) for the single task lasso bandit approach. We present here a sketch summarizing the key steps. Full technical details are provided in Appendix D.

Proof Sketch of Theorem 1. We consider the instantaneous multi-task regret at round n

$$\bar{R}_n = \sum_{t=1}^T R_{t,n} = \sum_{t=1}^T \langle \mathbf{x}_{t,n}^*, \mathbf{w}_t \rangle - \langle \mathbf{x}_{t,n}, \mathbf{w}_t \rangle.$$

During the first $N_0(\delta)$ rounds, the RSC condition may not be satisfied. During this phase, using a simple conditioning argument, we obtain the following bound. We have with probability at least $1 - \delta$, for any $n \in [N_0]$

$$\sum_{n=1}^{N_0} \bar{R}_n \lesssim T L N_0 \sqrt{\log(eTN_0 K \delta^{-1})}.$$

Starting from the $N_0 + 1$ round, we bound the instantaneous regret as follows:

$$\bar{R}_n \leq \sum_{n=N_0+1}^N \sum_{t=1}^T \langle \mathbf{x}_{t,n}^* - \mathbf{x}_{t,n}, \mathbf{w}_t - \widehat{\mathbf{w}}_{t,n} \rangle.$$

Using another conditioning, we get with probability at least $1 - \delta$, for any $t \in [T]$, $n \in [N]$,

$$\bar{R}_n \lesssim \sqrt{T} \left\| \mathbf{W} - \widehat{\mathbf{W}}_n \right\|_F \sqrt{\log(eTNK \delta^{-1})}.$$

This bound is of interest if we can guarantee that $\widehat{\mathbf{W}}_n$ is an accurate estimate of \mathbf{W} , meaning that $\|\mathbf{W} - \widehat{\mathbf{W}}_n\|_F$ is small.

To this end, we first prove that $\widehat{\Sigma}_n$ satisfies the RSC condition whp starting from round N_0 . The argument is based on concentration bounds for martingales combined with a novel geometric analysis of the cone of matrix C_r ; see Lemma 2.

Proposition 2. *Let Assumptions 1 and 2 be satisfied. Assume in addition that Σ satisfies the RSC condition. Then, for any $\delta \in (0, 1)$ and any $n \geq N_0(\delta)$, with probability at least $1 - \delta$, the multi-task empirical matrix $\widehat{\Sigma}_n$ satisfies the RSC condition with constant*

$$\kappa(\widehat{\Sigma}_n) \geq \frac{\kappa(\Sigma)}{4\nu\omega_{\mathcal{X}}} > 0.$$

Consequently, this means that we can now use Lemma 1 and Proposition 1 which guarantee with probability at least $1 - \delta$, simultaneously for any $n \in [N_0, N]$

$$\|\widehat{\mathbf{W}}_n - \mathbf{W}\|_F \lesssim \sqrt{\frac{r}{n}} \left(\sqrt{T + d + \log\left(\frac{4N(d+T)}{\delta}\right)} \cdot \log^{3/2}\left(\frac{8N(d+T)}{\delta}\right) \right).$$

Summing over n , we obtain with probability at least $1 - \delta$ up to logarithmic factors

$$\sum_{n=N_0+1}^N \bar{R}_s \leq C\sqrt{rT(T+d)N},$$

where $C = C\left(L, \eta, \sigma, C_{\mathbf{z}}, \kappa(\Sigma), \max_{1 \leq k \leq K} \{\|\Sigma_k\|_{\text{op}}^{1/2}\}\right)$ is a finite constant under our assumptions.

An union bound summing the regrets for the first phase $n \leq N_0$ and the second phase $n > N_0$ gives the result (up to a rescaling of the constants). \square

5.2 Result Discussion

We now discuss the implication of Theorem 1 and compare it to previous approaches to multi-task representation learning in the bandit setting.

Advantage over ITL. Notice that running any T independent policies (ITL) with the linear contextual bandit setting defined in Section 3.1 would yield at best a regret bound of order $T\sqrt{dN}$, up to logs; see e.g. (Lattimore and Szepesvári, 2020, Chapter 19.4, comment 5). Since this regret bound is always larger than the upper bound in Theorem 1 for the proposed MTL strategy, there is a gain in using our method. In particular, if $T > d$, discarding logarithmic factors, the bound for our method is smaller by a factor of order $O(\sqrt{d/r})$, while for $d > T$ the gain is of order $O(\sqrt{T/r})$.

Minimax Optimality. Theorem 2 in (Yang et al., 2020) provides a matching minimax lower bound to our Theorem 1 (up to logarithmic terms). In particular this implies that in the regime $T \geq d$, our policy achieves the minimax regret $T\sqrt{rN}$ (up to logs). This corresponds to the performance of the oracle policy which knows the true underlying representation a-priori.

Comparison to SOTA Approaches. Our upper bound compares favorably to Theorem 1 in (Yang et al., 2020) for the finitely many arms setting. Their result assumes Gaussian arms with non singular covariance matrices. They also assume that $\text{rank}(\mathbf{W}) = r$ with **known** r and $K, T \leq \text{poly}(d)$, $N \geq d^2$. Then Theorem 1 in (Yang et al., 2020) guarantees for their MLinGreedy policy that

$$\mathbb{E}[\bar{R}(T, N)] = O\left(\left(T\sqrt{rN} + \sqrt{rdTN} \right) \cdot \sqrt{\log(NKT) \log(NTdr) \log \log N} \right).$$

We stress out that their policy requires the knowledge of the rank r whereas our policy does not. Notably, their regret bound requires $N \geq d^2$ rounds to be valid whereas our regret bound is valid as soon as $N \gtrsim r^2$ (up to logs). Moreover, their analysis requires the invertibility of the arms covariance, whereas we only need the less restrictive RSC condition. Finally we also extend the result to sub-Gaussian arm distributions. With our notation, (Hu et al. (2021) obtained a regret bound of order $O(T\sqrt{drN} + d\sqrt{rTN})$ up to logs for their MTLR-OFUL policy in the infinitely many arms setting, provided the rank r is known to their policy.

6 EXPERIMENTS

In this section we validate the policy proposed in Section 5. The experiments displayed in Figures 1 and 2 compare 3 different policies: the Trace-Norm Bandit approach of Algorithm 1 for different choices of the regularization parameter, the Oracle Policy which knows the low dimensional representation $\mathbf{B} \in \mathbb{R}^{d \times r}$ to select the arm to play at each round and the ITL policy which solves each d -dimensional task separately. In order to compute the trace-norm estimator (Eq. 4) we adopt the accelerated gradient method proposed in (Ji and Ye, 2009). In all the experiments, we report results averaged over 5 repetitions.

Data Processing. We conduct numerical experiments on different configurations. Specifically, we analyze the impact of the number of tasks T , the dimension d , noise variance σ^2 and the choice of the regularization parameter. As metric we consider the cumulative reward averaged on all the tasks $\bar{R}(T, N)/T$ as a function of the number of rounds N . The arm set is generated randomly from a standard Gaussian distribution. The task matrix \mathbf{W} has been chosen to be of rank $r \ll d$ with randomly generated gaussian entries.

Result Discussion. Figures 1 and 2 indicate that the pro-

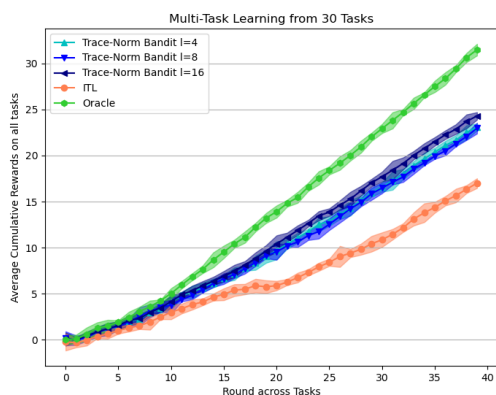
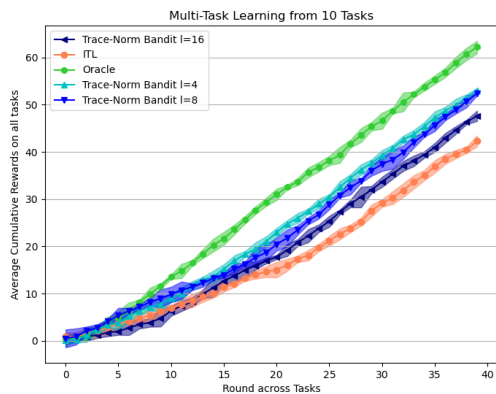


Figure 1: Averaged cumulative reward over all tasks for $T = 10$ (top) and $T = 30$ (bottom). Each task lasts for $N = 40$ rounds, has $K = 10$ arms with $d = 20$ features, noise variance $\sigma^2 = 1$.

posed multi-task approach performs favorably over independent task learning. Figure 1 highlights that our trace-norm policy performs well and significantly better than ITL policy as long as the regularization parameter used in Eq. (4) is taken large enough according to theory:

$$\lambda_n = l \left[\left(\frac{T+d}{n} + \frac{\log \frac{2}{\delta}}{n} \right) \vee \left(\sqrt{\frac{T+d}{n}} + \sqrt{\frac{\log \frac{2}{\delta}}{n}} \right) \right].$$

In addition, we also observe that the performance of the trace norm bandit improves and tends to that of the oracle policy as the number of tasks increases. In Figure 2, we investigate the impact of noise variance on the performance of the different policies. We used a higher value of dimension $d = 50$ as it is well-known that the impact of the noise becomes more problematic in high-dimension. We observe that the trace-norm policy is significantly less impacted by increased noise variance than the ITL policy. Indeed the performance of the ITL policy degrades by 48% as the noise variance increases from 1 to 9. Comparatively, the oracle policy and the trace-norm policy only incur a degradation of about 25% as the noise variance increases. This may be due to the fact

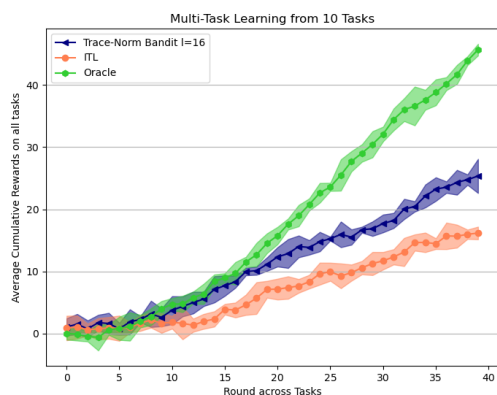
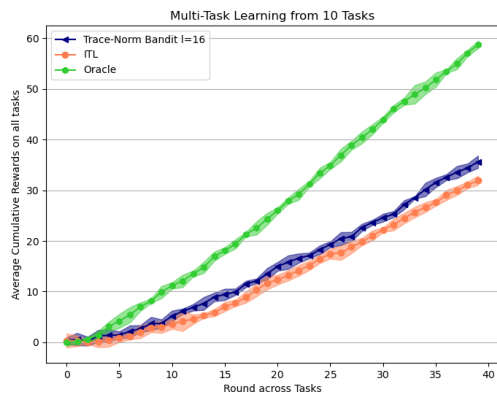


Figure 2: Averaged cumulative reward over all tasks. Each task lasts for $N = 40$ rounds, has $K = 10$ arms with $d = 50$ features and noise variance $\sigma^2 = 1$ (top), $\sigma^2 = 9$ (bottom).

that the oracle works in a subspace of dimension r , and that the trace-norm policy is able to perform dimensionality reduction through the nuclear norm regularization whereas the ITL policy works in the whole d -dimensional space.

7 LIMITATIONS AND FUTURE WORK

We have studied the benefit of multi-task representation learning in the setting of linear contextual bandit tasks. We proposed a novel bandit policy based on trace-norm regularization which is computationally efficient and does not need knowledge of the rank of the underlying task matrix. We derived an upper bound for the multi-task regret of the proposed policy, showing that it is effective in comparison to learning the tasks independently. Additionally, this regret bound is minimax optimal and, in the regime $T \geq d$, our policy's regret matches the one of the oracle policy which knows the low-rank common representation.

In this work we have restricted our analysis to the case of linear feature learning. Additionally, we still require the designed policy to know one problem parameter which is the variance associated to the noisy term. Relying on (Maurer

et al., 2016), in the future an interesting extension would be to go beyond the linearity of the shared representation. Secondly, inspired by (Belloni et al., 2011) a challenge one may try to solve is to design a fully parameter free policy. Namely, a policy which does not require to know the noise variance σ^2 . Finally an important future direction is to consider the meta-learning setting, in which the tasks are observed sequentially and the goal is to minimize the regret on future yet-unseen tasks.

Acknowledgements

This work was supported in part by the Chaire Business Analytic for Future Banking, PNRR MUR project PE0000013-FAIR, and the European Union (Project 101070617).

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *NIPS*, volume 11, pages 2312–2320, 2011.
- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9, 2012.
- Andreas Argyriou, Theodoros Evgeniou, and Massimiliano Pontil. Convex multi-task feature learning. *Machine learning*, 73(3):243–272, 2008.
- Kaito Ariu, Kenshi Abe, and Alexandre Proutière. Thresholded lasso bandit. *arXiv preprint arXiv:2010.11994*, 2020.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Baruch Awerbuch and Robert Kleinberg. Online linear optimization and adaptive routing. *Journal of Computer and System Sciences*, 74(1):97–114, 2008.
- Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294, 2020.
- Hamsa Bastani, Mohsen Bayati, and Khashayar Khosravi. Mostly exploration-free algorithms for contextual bandits. *Management Science*, 67(3):1329–1349, 2021.
- Soumya Basu, Branislav Kveton, Manzil Zaheer, and Csaba Szepesvári. No regrets for learning the prior in bandits. *arXiv preprint arXiv:2107.06196*, 2021.
- Jonathan Baxter. A model of inductive bias learning. *Journal of artificial intelligence research*, 12:149–198, 2000.
- Alexandre Belloni, Victor Chernozhukov, and Lie Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.
- Peter Bühlmann and Sara Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- Daniele Calandriello, Alessandro Lazaric, and Marcello Restelli. Sparse multi-task reinforcement learning. In *NIPS*, 2014.
- Leonardo Cella and Massimiliano Pontil. Multi-task and meta-learning with sparse linear bandits. In *The Conference on Uncertainty in Artificial Intelligence*, 2021.
- Leonardo Cella, Alessandro Lazaric, and Massimiliano Pontil. Meta-learning with stochastic linear bandits. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1360–1370. PMLR, 13–18 Jul 2020.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- Giulia Denevi, Carlo Ciliberto, Dimitris Stamos, and Massimiliano Pontil. Incremental learning-to-learn with statistical guarantees. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence*, pages 457–466, 2018.
- Giulia Denevi, Dimitris Stamos, Carlo Ciliberto, and Massimiliano Pontil. Online-within-online meta-learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Carlo D’Eramo, Davide Tateo, Andrea Bonarini, Marcello Restelli, and Jan Peters. Sharing knowledge in multi-task deep reinforcement learning. In *International Conference on Learning Representations*, 2019.
- Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *NIPS*, volume 23, pages 586–594, 2010.
- Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR, 2020.
- Botao Hao, Tor Lattimore, and Mengdi Wang. High-dimensional sparse linear bandits. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10753–10763. Curran Associates, Inc., 2020.
- Jiachen Hu, Xiaoyu Chen, Chi Jin, Lihong Li, and Liwei Wang. Near-optimal representation learning for linear bandits and linear rl. In *International Conference on Machine Learning*, pages 4349–4358. PMLR, 2021.
- Shuiwang Ji and Jieping Ye. An accelerated gradient method for trace norm minimization. In *International Conference on Machine Learning*, pages 457–464, 2009.

- Gi-Soo Kim and Myunghee Cho Paik. Doubly-robust lasso bandit. *Advances in Neural Information Processing Systems*, 32:5877–5887, 2019.
- V. Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems: École D’Été de Probabilités de Saint-Flour XXXVIII-2008*. Lecture Notes in Mathematics. Springer, 2011. ISBN 9783642221460. URL https://books.google.fr/books?id=D5Jxen3_xkAC.
- Vladimir Koltchinskii, Karim Lounici, and Alexandre B Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.
- Ilja Kuzborskij, Leonardo Cella, and Nicolò Cesa-Bianchi. Efficient linear bandits through matrix sketching. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 177–185. PMLR, 2019.
- Branislav Kveton, Csaba Szepesvári, Anup Rao, Zheng Wen, Yasin Abbasi-Yadkori, and S Muthukrishnan. Stochastic low-rank bandits. *arXiv preprint arXiv:1712.04644*, 2017.
- Branislav Kveton, Mikhail Konobeev, Manzil Zaheer, Chiwei Hsu, Martin Mladenov, Craig Boutilier, and Csaba Szepesvari. Meta-thompson sampling. *arXiv preprint arXiv:2102.06129*, 2021.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning*, pages 2071–2080, 2017.
- Karim Lounici, Massimiliano Pontil, Sara Van De Geer, and Alexandre B Tsybakov. Oracle inequalities and optimal inference under group sparsity. *The annals of statistics*, 39(4):2164–2204, 2011.
- Yangyi Lu, Amirhossein Meisami, and Ambuj Tewari. Low-rank generalized linear bandit problems. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 460–468. PMLR, 13–15 Apr 2021.
- Andreas Maurer and Massimiliano Pontil. Excess risk bounds for multitask learning with trace norm regularization. In *Conference on Learning Theory*, pages 55–76. PMLR, 2013.
- Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.
- Sahand Negahban and Martin J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, 39(2):1069–1097, 2011.
- Min-hwan Oh, Garud Iyengar, and Assaf Zeevi. Sparsity-agnostic lasso bandit. In *Proceedings of the 38-th International Conference on Machine Learning*, volume 139, pages 8271–8280, 2021.
- Paat Rusmevichientong and John N. Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, page 395–411, 2010.
- Weiwei Shen, Jun Wang, Yu-Gang Jiang, and Hongyuan Zha. Portfolio choices with orthogonal bandit learning. In *Twenty-fourth international joint conference on artificial intelligence*, 2015.
- Max Simchowitz, Christopher Tosh, Akshay Krishnamurthy, Daniel Hsu, Thodoris Lykouris, Miroslav Dudík, and Robert E Schapire. Bayesian decision-making under misspecified priors with applications to meta-learning. *arXiv preprint arXiv:2107.01509*, 2021.
- Nilesh Tripuraneni, Chi Jin, and Michael Jordan. Provable meta-learning of linear representations. In *International Conference on Machine Learning*, pages 10434–10443. PMLR, 2021.
- Joel Tropp. Freedman’s inequality for matrix martingales. *Electronic Communications in Probability*, 16(none):262 – 270, 2011. doi: 10.1214/ECP.v16-1624. URL <https://doi.org/10.1214/ECP.v16-1624>.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018. doi: 10.1017/9781108231596.
- Xue Wang, Mingcheng Wei, and Tao Yao. Minimax concave penalized multi-armed bandit model with high-dimensional covariates. In *International Conference on Machine Learning*, pages 5200–5208, 2018.
- Michael Woodroffe. A one-armed bandit problem with a concomitant variable. *Journal of the American Statistical Association*, 74(368):799–806, 1979.
- Jiaqi Yang, Wei Hu, Jason D Lee, and Simon Shaolei Du. Impact of representation learning in linear bandits. In *International Conference on Learning Representations*, 2020.

APPENDIX

This appendix provides full proofs of the results stated in the main body of the paper. It is organized as follows:

- Appendix A contains the proof of Lemma 1, which is the oracle inequality associated to the error $\Delta_{n+1} = \widehat{\mathbf{W}}_{n+1} - \mathbf{W}$ considering non i.i.d. data.
- Appendix B presents the proof of Proposition 1 on the control to the operator norm of matrix $\mathbf{D}_n = \sum_{t=1}^T \sum_{i=1}^n \eta_{t,i} \mathbf{x}_{t,i} \otimes \mathbf{e}_t$.
- Appendix C contains the proof of Proposition 2. It relies on a matrix perturbation argument and a novel analysis of the uniform deviation of the empirical arms covariance on the cone \mathcal{C}_r of approximately low-rank matrices (Lemmas 2 and 3). These results are then combined to relate the RSC constant $\kappa(\overline{\Sigma})$ of $\overline{\Sigma}$ to that of its empirical counterpart $\kappa(\widehat{\Sigma}_n)$.
- Appendix D contains the proof for the regret upper bound associated to the policy of Alg. 1.
- In Appendix E, we provide additional numerical experiments.

A PROOF OF LEMMA 1

By definition of $\widehat{\mathbf{W}}_n$, we have for any $\mathbf{W} \in \mathbb{R}^{d \times T}$,

$$\frac{1}{n} \sum_{t=1}^T \left\| \mathbf{y}_{t,n} - \mathbf{X}_{t,n} [\widehat{\mathbf{W}}_{n+1}]_t \right\|_2^2 + \lambda_n \left\| \widehat{\mathbf{W}}_{n+1} \right\|_* \leq \frac{1}{n} \sum_{t=1}^T \left\| \mathbf{y}_{t,n} - \mathbf{X}_{t,n} [\mathbf{W}]_t \right\|_2^2 + \lambda_n \left\| \mathbf{W} \right\|_*.$$

We define the error matrix $\Delta_{n+1} = \widehat{\mathbf{W}}_{n+1} - \mathbf{W}$ and introduce the following operator $\mathcal{A} : \mathbb{R}^{d \times T} \rightarrow \mathbb{R}^{n \times T}$ and its adjoint $\mathcal{A}^* : \mathbb{R}^{n \times T} \rightarrow \mathbb{R}^{d \times T}$ as

$$\begin{aligned} [\mathcal{A}(\mathbf{W})]_n^t &= \langle \mathbf{x}_{t,n} \mathbf{e}_t^\top, \mathbf{W} \rangle = \text{Tr}(\mathbf{W} \mathbf{e}_t \mathbf{x}_{t,n}) = \mathbf{x}_{t,n}^\top \mathbf{w}_t \\ \mathcal{A}^*(\mathbf{H}_n) &= \sum_{t=1}^T \sum_{i=1}^n \mathbf{x}_{t,i} \mathbf{e}_i^\top \eta_{t,i} = \mathbf{D}_n \in \mathbb{R}^{d \times T}, \end{aligned}$$

where $\boldsymbol{\eta}_{t,n} = (\eta_{t,1}, \dots, \eta_{t,n})^\top \in \mathbb{R}^n$ and $\mathbf{H}_n := (\boldsymbol{\eta}_{1,n}, \dots, \boldsymbol{\eta}_{T,n}) \in \mathbb{R}^{n \times T}$.

Using this notation the following hold

$$\begin{aligned} \frac{1}{n} \|\mathcal{A}(\Delta_{n+1})\|_{\text{F}}^2 &= \frac{1}{n} \sum_{t=1}^T \left\| [\mathcal{A}(\Delta_{n+1})]_t \right\|_2^2 \leq \frac{1}{n} \langle \mathbf{H}_n, \mathcal{A}(\Delta_{n+1}) \rangle + \lambda_n \left(\|\mathbf{W}\|_* - \left\| \widehat{\mathbf{W}}_{n+1} \right\|_* \right) \\ &\leq \frac{1}{n} \langle \mathbf{H}_n, \mathcal{A}(\Delta_{n+1}) \rangle + \lambda_n \left(\left\| \widehat{\mathbf{W}}_{n+1} + \Delta_{n+1} \right\|_* - \left\| \widehat{\mathbf{W}}_{n+1} \right\|_* \right) \\ &\leq \frac{1}{n} \langle \mathbf{H}_n, \mathcal{A}(\Delta_{n+1}) \rangle + \lambda_n \left(\|\Delta_{n+1}\|_* \right). \end{aligned}$$

Considering now the first term on the RHS and applying Holder's inequality we have

$$\frac{1}{n} |\langle \mathbf{H}_n, \mathcal{A}(\Delta_{n+1}) \rangle| = \frac{1}{n} |\langle \mathcal{A}^*(\mathbf{H}_n), \Delta_{n+1} \rangle| \leq \frac{1}{n} \|\mathcal{A}^*(\mathbf{H}_n)\|_{\text{op}} \|\Delta_{n+1}\|_*.$$

Now, considering $\lambda_n \geq \frac{1}{n} \|\mathcal{A}^*(\mathbf{H}_n)\|_{\text{op}}$ the following holds:

$$\frac{1}{n} \|\mathcal{A}(\Delta_{n+1})\|_{\text{F}}^2 \leq 2\lambda_n \|\Delta_{n+1}\|_*. \quad (13)$$

Relying on (Negahban and Wainwright, 2011, Lemma 1) we can decompose the error matrix Δ_{n+1} as $\Delta'_{n+1} + \Delta''_{n+1}$ such that Δ'_{n+1} is of rank at most $2r$ and

$$\|\Delta_{n+1}\|_* \leq 4 \|\Delta'_{n+1}\|_*. \quad (14)$$

Assuming the RSC condition to be met with constant $\kappa(\widehat{\boldsymbol{\Sigma}}_n)$, starting from equation (13) we get

$$\|\Delta_{n+1}\|_F^2 \leq \frac{\|\mathcal{A}(\Delta_{n+1})\|_F^2}{2n\kappa(\widehat{\boldsymbol{\Sigma}}_n)} \leq \frac{\lambda_n \|\Delta_{n+1}\|_*}{\kappa(\widehat{\boldsymbol{\Sigma}}_n)}.$$

These last two results combined with (Negahban and Wainwright, 2011, Lemma 1) give that $\|\Delta'_{n+1}\|_* \leq \sqrt{2r} \|\Delta'_{n+1}\|_F$, from which we conclude that

$$\|\Delta_{n+1}\|_F \leq \frac{32\lambda_n\sqrt{r}}{\kappa(\widehat{\boldsymbol{\Sigma}}_n)}.$$

Symbol	Description
$[n]$	The set $\{1, \dots, n\}$, given a positive integer n
T	Number of tasks
N	Time horizon associated to each single task
d	Dimension of context vectors
$a \vee b$	The maximum between a and b ($\max(a, b)$)
$\ \cdot\ _{\psi_2}$	The sub-Gaussian norm with $\psi_2(s) = e^{s^2} - 1$ (See e.g. page 215 in Koltchinskii (2011))
$\mathbf{e}_1, \dots, \mathbf{e}_T \in \mathbb{R}^T$	The standard basis indicator vectors, i.e. $\mathbf{e}_{t,j} = 1$, if $j = t$ and 0 otherwise, for all $t, j \in [T]$
$\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_T] \in \mathbb{R}^{d \times T}$	Matrix of T regression tasks (we also use the notation $[\mathbf{W}]_t \equiv \mathbf{w}_t$, $t \in [T]$)
r	Rank of the task matrix \mathbf{W}
K	Number of arms
p	Joint distribution on \mathbb{R}^{dK} (from which K arm vectors are sample)
$\mathcal{D}_{t,n}$, $t \in [T]$, $n \in [N]$	Decision sets (each containing K arm vectors) sampled i.i.d. from p
$\mathbf{x}_{t,n} \in \mathcal{D}_{t,n}$	Arm vector chosen in task $t \in [T]$ at round $n \in [N]$
$\mathbf{x}_{t,n}^* \in \mathcal{D}_{t,n}$	Optimal arm vector in task $t \in [T]$ during round $n \in [N]$
$\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$	Theoretical covariance matrix - see eq. (5)
$\boldsymbol{\Sigma}_{t,n} \in \mathbb{R}^{d \times d}$	Adapted covariance matrix for task t at round n ; see eq. (32)
$\widehat{\boldsymbol{\Sigma}}_{t,n} \in \mathbb{R}^{d \times d}$	Empirical covariance matrix for task t at round n ; see eq. (6)
$\overline{\boldsymbol{\Sigma}} \in \mathbb{R}^{dT \times dT}$	T -block diagonal matrix $\text{diag}(\boldsymbol{\Sigma}, \dots, \boldsymbol{\Sigma})$
$\overline{\boldsymbol{\Sigma}}_n \in \mathbb{R}^{dT \times dT}$	T -block diagonal matrix $\text{diag}(\boldsymbol{\Sigma}_{1,n}, \dots, \boldsymbol{\Sigma}_{T,n})$
$\widehat{\overline{\boldsymbol{\Sigma}}}_n \in \mathbb{R}^{dT \times dT}$	T -block diagonal matrix $\text{diag}(\widehat{\boldsymbol{\Sigma}}_{1,n}, \dots, \widehat{\boldsymbol{\Sigma}}_{T,n})$
$\ \mathbf{x}\ $, $\ \mathbf{x}\ _1$, $\ \mathbf{x}\ _\infty$	Euclidean, ℓ_1 and maximum norm associated to a vector \mathbf{x}
$[\mathbf{A}]_t$	The t -th column of matrix $\mathbf{A} \in \mathbb{R}^{d \times T}$
$\lambda_{\min}(\mathbf{A})$, $\lambda_{\max}(\mathbf{A})$	Minimum and maximum eigenvalues of a square symmetric matrix \mathbf{A}
$\sigma_{\min}(\mathbf{A})$, $\sigma_{\max}(\mathbf{A})$	Minimum and maximum singular values of matrix \mathbf{A}
$\ \mathbf{A}\ _*$	Trace norm of matrix \mathbf{A} (sum of its singular values)
$\ \mathbf{A}\ _F$	Frobenius norm of matrix \mathbf{A} (ℓ_2 norm of matrix elements / singular values)
$\ \mathbf{A}\ _{\text{op}}$	Operator norm of matrix \mathbf{A} (maximum singular value)

B PROOF OF PROPOSITION 1

B.1 Preliminary results

We consider the stochastic process $\{\mathbf{M}_n\}_{n \geq 0}$ defined as $\mathbf{M}_0 = 0$ a.s. and for any $n \geq 1$

$$\mathbf{M}_n = \sum_{t=1}^T \mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} - \mathbb{E}[\mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} | \bar{\mathcal{F}}_{n-1}].$$

By definition of the Trace-Norm bandit (Algorithm 1), $\{\mathbf{M}_n\}_{n \geq 0}$ is a $\bar{\mathcal{F}}_{n-1}$ -martingale. Furthermore, given the past history $\bar{\mathcal{F}}_{n-1}$, we select at round n for each task t the arm $\mathbf{x}_{t,n} \in \mathcal{D}_{t,n}$ where the sets $\mathcal{D}_{t,n}$, $t \in [T]$, are mutually independent. This means that the arms $\mathbf{x}_{t,n}$, $t \in [T]$, are mutually independent given the past history $\bar{\mathcal{F}}_{n-1}$.

Next, we use a standard argument to control the operator norm of \mathbf{M}_n . Fix $\epsilon \in (0, 1/2)$. An ϵ -net \mathcal{N}_ϵ of \mathbb{S}^d is a subset of \mathbb{S}^d such that for any $\mathbf{u} \in \mathbb{S}^d$, there exists $\mathbf{v} \in \mathcal{N}_\epsilon$ such that $\|\mathbf{u} - \mathbf{v}\| \leq \epsilon$. Corollary 4.2.13 in Vershynin (2018) guarantees the existence of an ϵ -net \mathcal{N}_ϵ of \mathbb{S}^d such that

$$|\mathcal{N}_\epsilon| \leq \left(1 + \frac{2}{\epsilon}\right)^d. \quad (15)$$

Similarly to the first step of the proof of Theorem 4.4.5 in Vershynin (2018), we get

$$\|\mathbf{M}_n\|_{\text{op}} \leq \frac{1}{1-2\epsilon} \max_{\mathbf{u} \in \mathcal{N}_\epsilon} \{\langle \mathbf{M}_n \mathbf{u}, \mathbf{u} \rangle\} \leq \frac{1}{1-2\epsilon} \max_{\mathbf{u} \in \mathcal{N}_\epsilon} \left\{ \sum_{t=1}^T \langle \mathbf{x}_{t,n}, \mathbf{u} \rangle^2 - \mathbb{E}[\langle \mathbf{x}_{t,n}, \mathbf{u} \rangle^2 | \bar{\mathcal{F}}_{n-1}] \right\}. \quad (16)$$

In view of Assumption 1, we have $\langle \mathbf{x}_{t,n}, \mathbf{u} \rangle = \langle \mathbf{z}_{t,n}, \boldsymbol{\Sigma}_k^{1/2} \mathbf{u} \rangle$ and $\mathbb{E}[\langle \mathbf{x}_{t,n}, \mathbf{u} \rangle^2 | \bar{\mathcal{F}}_{n-1}] = \langle \boldsymbol{\Sigma}_{k(t)} \mathbf{u}, \mathbf{u} \rangle$ for some $k(t) \in [K]$. We apply now the Hanson-Wright's inequality conditionally on $\bar{\mathcal{F}}_{n-1}$ to get for any $x > 0$

$$\mathbb{P} \left(\sum_{t=1}^T \langle \mathbf{z}_{t,n}, \boldsymbol{\Sigma}_{k(t)}^{1/2} \mathbf{u} \rangle^2 - \langle \boldsymbol{\Sigma}_{k(t)} \mathbf{u}, \mathbf{u} \rangle \geq C \left(\sqrt{T \max_{k \in [K]} \left\{ \left\| \boldsymbol{\Sigma}_k^{1/2} \mathbf{u} \otimes \mathbf{u} \boldsymbol{\Sigma}_k^{1/2} \right\|_F^2 \right\}} x + \max_{k \in [K]} \{ \|\boldsymbol{\Sigma}_k \mathbf{u} \otimes \mathbf{u} \boldsymbol{\Sigma}_k\|_{\text{op}} \} x \right) | \bar{\mathcal{F}}_{n-1} \right) \leq e^{-x},$$

where $C > 0$ is a numerical constant which can depend only on $C_{\mathbf{z}}$. Note that $\left\| \boldsymbol{\Sigma}_k^{1/2} \mathbf{u} \otimes \mathbf{u} \boldsymbol{\Sigma}_k^{1/2} \right\|_F = \left\| \boldsymbol{\Sigma}_k^{1/2} \mathbf{u} \otimes \mathbf{u} \boldsymbol{\Sigma}_k^{1/2} \right\|_{\text{op}} \|\mathbf{u}\|^2 \leq \|\boldsymbol{\Sigma}_k\|_{\text{op}}$. Hence we get for any $x > 0$

$$\mathbb{P} \left(\sum_{t=1}^T \langle \mathbf{z}_{t,n}, \boldsymbol{\Sigma}_{k(t)} \mathbf{u} \rangle^2 - \langle \boldsymbol{\Sigma}_{k(t)} \mathbf{u}, \mathbf{u} \rangle \geq C \max_{k \in [K]} \{ \|\boldsymbol{\Sigma}_k\|_{\text{op}} \} (\sqrt{T} x + x) | \bar{\mathcal{F}}_{n-1} \right) \leq e^{-x}.$$

We define now the event

$$\Omega_n = \bigcap_{\mathbf{u} \in \mathcal{N}_\epsilon} \left\{ |\langle \mathbf{M}_n \mathbf{v}, \mathbf{u} \rangle| \leq C \max_{k \in [K]} \{ \|\boldsymbol{\Sigma}_k\|_{\text{op}} \} (\sqrt{T} x + x) \right\}.$$

Set $\epsilon = 1/4$. A simple union bound combining (15) with the last two displays gives

$$\mathbb{P}(\Omega_n^c) \leq |\mathcal{N}_\epsilon| e^{-x} \leq 2e^{d \log(9) - x}.$$

Now we set $x = \log(4\delta^{-1}) + d \log(9)$ for some $\delta \in (0, 1)$. Consequently, we obtain that

$$\mathbb{P}(\Omega_n) \geq 1 - \delta/4.$$

It follows, in view of (37), with probability at least $1 - \delta$

$$\|\mathbf{M}_n\|_{\text{op}} \leq C' \max_{k \in [K]} \{ \|\boldsymbol{\Sigma}_k\|_{\text{op}} \} \left(\sqrt{T(d + \log(4\delta^{-1}))} + d + \log(4\delta^{-1}) \right), \quad (17)$$

for some numerical constant $C' > 0$ that can depend only on $C_{\mathbf{z}}$.

We note that

$$\left\| \sum_{t=1}^T \mathbb{E} [\mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} | \bar{\mathcal{F}}_{n-1}] \right\|_{\text{op}} \leq T \max_k \{\|\boldsymbol{\Sigma}_k\|_{\text{op}}\},$$

since $\sum_{t=1}^T \mathbb{E} [\mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} | \bar{\mathcal{F}}_{n-1}] = \sum_{t=1}^T \boldsymbol{\Sigma}_{k(t)}$, where $k(t) \in [K]$. Combining this observation with (17), we prove that the following event

$$\Omega_n = \left\{ \left\| \sum_{t=1}^T \mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} \right\|_{\text{op}} \leq \max_{k \in [K]} \{\|\boldsymbol{\Sigma}_k\|_{\text{op}}\} \left(T + C' \left(\sqrt{T(d + \log(4N\delta^{-1}))} + d + \log(4N\delta^{-1}) \right) \right) \right\} \quad (18)$$

satisfies $\mathbb{P}(\Omega_n) \geq 1 - \delta/(4N)$. Next we also introduce the event

$$\bar{\Omega}_n = \bigcap_{l=0}^n \Omega_l \in \bar{\mathcal{F}}_n, \quad (19)$$

with Ω_0 being the whole sample space. The Bayes rule gives

$$\mathbb{P}(\bar{\Omega}_N) = \prod_{n=1}^N \mathbb{P} \left(\Omega_n \mid \bigcap_{k=0}^{n-1} \Omega_k \right) \geq \left(1 - \frac{\delta}{4N} \right)^N.$$

Bernoulli's inequality ($(1+x)^n \geq 1+nx$ for any $x > -1$ and integer $n \geq 1$) gives

$$\mathbb{P}(\bar{\Omega}_N) = \prod_{n=1}^N \mathbb{P} \left(\Omega_n \mid \bigcap_{k=0}^{n-1} \Omega_k \right) \geq 1 - \frac{\delta}{4}. \quad (20)$$

Define for any $n \in [N]$ the events

$$\Omega'_n = \bigcap_{t=1}^T \left\{ |\eta_{t,n}| \leq c_\eta \sqrt{\log(8TN\delta^{-1})} \right\}, \quad (21)$$

and

$$\bar{\Omega}'_n = \bigcap_{l=0}^n \Omega'_l, \quad (22)$$

with Ω'_0 being the whole sample space.

Assumption 4 guarantees that for any $\delta \in (0, 1)$, with probability at least $1 - \delta$

$$\mathbb{P}(\Omega'_n | \bar{\mathcal{F}}_{n-1}) \geq 1 - \frac{\delta}{4N}. \quad (23)$$

and

$$\mathbb{P}(\bar{\Omega}'_N | \bar{\mathcal{F}}_{N-1}) \geq 1 - \frac{\delta}{4}. \quad (24)$$

B.2 Main proof

Checking the martingale structure. Let us define the stochastic process $(\mathbf{D}_n)_{n \geq 0}$ as $\mathbf{D}_0 = 0$ a.s. and

$$\mathbf{D}_n := \sum_{t=1}^T \sum_{i=1}^n \eta_{t,i} \mathbf{x}_{t,i} \mathbf{e}_t^\top \in \mathbb{R}^{d \times T}. \quad (25)$$

By definition of the Trace-Norm bandit (Algorithm 1), given the past history $\bar{\mathcal{F}}_{n-1}$, we select at round n the arms for the T tasks: $\mathbf{x}_{t,n} \in \mathcal{D}_{t,n}$ with $\mathcal{D}_{t,n}$ independent of $(\eta_{t,n})_{1 \leq t \leq T}$. This means that

$$(\mathbf{x}_{t,n})_{1 \leq t \leq T} \perp\!\!\!\perp (\eta_{t,n})_{1 \leq t \leq T} \Big| \bar{\mathcal{F}}_{n-1}.$$

Consequently, under Assumption 1 and 4, $(\mathbf{D}_n)_{n \geq 0}$ is a square-root integrable martingale adapted to the filtration $\{\overline{\mathcal{F}}_n\}_{n \geq 0}$. We would like to apply the Freedman inequality for matrix martingales (Tropp, 2011, Corollary 1.3). However this result is for bounded martingales. Therefore it is not directly applicable to $(\mathbf{D}_n)_{n \geq 0}$. To remedy this difficulty. We introduce the following stopping times:

$$\tau_1 = \inf \left\{ n \geq 0 : \left\| \sum_{t=1}^T \mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} \right\|_{\text{op}} \geq \max_{k \in [K]} \{\|\boldsymbol{\Sigma}_k\|_{\text{op}}\} \left(T + C' \left(\sqrt{T(d + \log(4N\delta^{-1}))} + d + \log(4N\delta^{-1}) \right) \right) \right\}, \quad (26)$$

and

$$\tau_2 = \inf \left\{ n \geq 0 : \min_{t \in [T]} \{|\eta_{t,n}|\} \geq c_\eta \sigma \sqrt{\log(8TN\delta^{-1})} \right\}. \quad (27)$$

Again by definition of the Trace-Norm bandit and assumptions on the noise, τ_1 and τ_2 are both stopping times relative to the filtration $\{\overline{\mathcal{F}}_n\}_{n \geq 0}$, so is $\tau = \tau_1 \wedge \tau_2$. Hence the stopped process $\{\mathbf{D}_n^\tau\}_{n \geq 0}$ defined as $\mathbf{D}_n^\tau = \mathbf{D}_{n \wedge \tau}$ is also a martingale adapted to the filtration $\{\overline{\mathcal{F}}_n\}_{n \geq 0}$. Furthermore $\{\mathbf{D}_n^\tau\}_{n \geq 0}$ is a bounded martingale. Hence we can apply the Freedman inequality to it.

By definition of τ_1 and (19), we have $\{\tau_1 > N\} = \overline{\Omega}_N$. Similarly for τ_2 in view of (22), we have $\{\tau_2 > N\} = \overline{\Omega}'_N$. Hence

$$\begin{aligned} \mathbb{P} \left(\|\mathbf{D}_n\|_{\text{op}} \geq t \right) &= \mathbb{P} \left(\{\mathbf{D}_n \geq t\} \cap \{\tau > N\} \right) + \mathbb{P} \left(\{\|\mathbf{D}_n\|_{\text{op}} \geq t\} \cap \{\tau \leq N\} \right) \\ &\leq \mathbb{P} \left(\|\mathbf{D}_n^\tau\|_{\text{op}} \geq t \right) + \mathbb{P} \left(\{\tau_1 \leq N\} \cup \{\tau_2 \leq N\} \right) \\ &= \mathbb{P} \left(\|\mathbf{D}_n^\tau\|_{\text{op}} \geq t \right) + \mathbb{P}(\tau_1 \leq N) + \mathbb{P}(\tau_2 \leq N) \\ &\leq \mathbb{P} \left(\|\mathbf{D}_n^\tau\|_{\text{op}} \geq t \right) + \delta/2, \end{aligned} \quad (28)$$

where we have used (20) and (24) in the last line.

Application of Freedman's inequality. We now use (Tropp, 2011, Corollary 1.3) to control $\{\mathbf{D}_n^\tau\}_{n \geq 0}$.

Theorem 2 (Corollary 1.3 in Tropp (2011)). *Consider a matrix martingale $\{\mathbf{Y}_n : n = 0, 1, 2, \dots\}$ whose values are matrices with dimension $d_1 \times d_2$, and let $\{\mathbf{X}_k : k = 1, 2, 3, \dots\}$ be the difference sequence. Assume that the difference sequence is uniformly bounded:*

$$\|\mathbf{X}_n\|_{\text{op}} \leq R \quad \text{almost surely for } n = 1, 2, 3, \dots$$

Define two predictable quadratic variation processes for this martingale:

$$\begin{aligned} \mathbf{W}_{\text{col},k} &:= \sum_{j=1}^n \mathbb{E}_{j-1} (\mathbf{X}_j \mathbf{X}_j^*) \quad \text{and} \\ \mathbf{W}_{\text{row},k} &:= \sum_{j=1}^n \mathbb{E}_{j-1} (\mathbf{X}_j^* \mathbf{X}_j) \quad \text{for } n = 1, 2, 3, \dots \end{aligned}$$

Then, for all $t \geq 0$ and $\overline{\sigma}^2 > 0$,

$$\mathbb{P} \left(\exists n \geq 0 : \|\mathbf{Y}_n\| \geq t \text{ and } \max\{\|\mathbf{W}_{\text{col},n}\|_{\text{op}}, \|\mathbf{W}_{\text{row},n}\|_{\text{op}}\} \leq \overline{\sigma}^2 \right) \leq (d_1 + d_2) \cdot \exp \left\{ -\frac{t^2/2}{\overline{\sigma}^2 + Rt/3} \right\}.$$

We now check the conditions of Theorem 2. In view of (18)-(19) and (21)-(22), we have on the event $\{\tau > N\} = \{\tau_1 >$

$N\} \cap \{\tau_2 > N\}$, for any $n \in [N]$,

$$\begin{aligned}
 \left\| \sum_{t=1}^T \eta_{t,n} \mathbf{x}_{t,n} \otimes \mathbf{e}_t \right\|_{\text{op}} &= \sqrt{\left\| \left(\sum_{t=1}^T \eta_{t,n} \mathbf{x}_{t,n} \otimes \mathbf{e}_t \right) \left(\sum_{t=1}^T \eta_{t,n} \mathbf{e}_t \otimes \mathbf{x}_{t,n} \right) \right\|_{\text{op}}} \\
 &= \sqrt{\left\| \sum_{t=1}^T \eta_{t,n}^2 \mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} \right\|_{\text{op}}} \leq \max_{t \in [T]} \{|\eta_{t,n}|\} \sqrt{\left\| \sum_{t=1}^T \mathbf{x}_{t,n} \otimes \mathbf{x}_{t,n} \right\|_{\text{op}}} \\
 &\leq c_\eta \sqrt{\log(8TN\delta^{-1})} \sqrt{\max_{k \in [K]} \{\|\Sigma_k\|_{\text{op}}\}} \left(T + C' \left(\sqrt{T(d + \log(4N\delta^{-1}))} + d + \log(4N\delta^{-1}) \right) \right) \\
 &\leq c_\eta \sqrt{1 + C'/2} \sqrt{\log(8TN\delta^{-1})} \sqrt{\max_{k \in [K]} \{\|\Sigma_k\|_{\text{op}}\}} (T + d + \log(4N\delta^{-1})).
 \end{aligned}$$

Next we have

$$\mathbf{W}_{\text{col}} = \sum_{s=1}^n \mathbb{E} \left[\sum_{t=1}^T \sum_{s=1}^n \eta_{t,s}^2 \mathbf{x}_{t,s} \otimes \mathbf{x}_{t,s} | \bar{\mathcal{F}}_{s-1} \right] = \sigma^2 \sum_{t=1}^T \sum_{s=1}^n \Sigma_{k(t,s)},$$

where $k(t, s) \in [K]$ for any t, s . Hence

$$\|\mathbf{W}_{\text{col}}\|_{\text{op}} \leq Tn\sigma^2 \max_{k \in [K]} \{\|\Sigma_k\|_{\text{op}}\}.$$

We proceed similarly for \mathbf{W}_{row}

$$\mathbf{W}_{\text{row}} = \sum_{s=1}^n \mathbb{E} \left[\sum_{t=1}^T \sum_{s=1}^n \eta_{t,s}^2 \|\mathbf{x}_{t,s}\|^2 \mathbf{e}_t \otimes \mathbf{e}_t | \bar{\mathcal{F}}_{s-1} \right] \leq n\sigma^2 \max_{k \in [K]} \{\text{tr}(\Sigma_k)\} I_T.$$

Hence, we get

$$\|\mathbf{W}_{\text{row}}\|_{\text{op}} \leq n\sigma^2 \max_{k \in [K]} \{\text{tr}(\Sigma_k)\} \leq nd\sigma^2 \max_{k \in [K]} \{\|\Sigma_k\|_{\text{op}}\}.$$

Define the event

$$A_n = \left\{ \frac{\|\mathbf{D}_n\|_{\text{op}}}{n} \leq \max_{k \in [K]} \{\|\Sigma_k\|_{\text{op}}^{1/2}\} \left(\sigma \sqrt{\frac{(d+T)(\log(2\delta^{-1}N(d+T)))}{n}} \right) \right. \\
 \left. \sqrt{\frac{c_\eta \sqrt{1 + C'/2} \sqrt{\log(8TN\delta^{-1})} \sqrt{(T + d + \log(4N\delta^{-1})) (\log(2\delta^{-1}N(d+T)))}}{n}} \right\}.$$

Applying Theorem 2, we get for any $t > 0$

$$\mathbb{P}(A_n^c | \bar{\Omega}_n) \leq \frac{\delta}{2N}. \tag{29}$$

An elementary argument combining the previous display with (55) gives

$$\mathbb{P}(A_n^c) = \mathbb{P}(A_n^c \cap \bar{\Omega}_n) + \mathbb{P}(A_n^c \cap \bar{\Omega}_n^c) \leq \mathbb{P}(A_n^c | \bar{\Omega}_n) + \mathbb{P}(\bar{\Omega}_n^c) \leq \frac{\delta}{N}.$$

From the previous display and an union bound, we immediately deduce, with probability at least $1 - \delta$,

$$\max_{1 \leq n \leq N} \frac{\|\mathbf{D}_n\|_{\text{op}}}{n} \leq \max_{k \in [K]} \{\|\Sigma_k\|_{\text{op}}^{1/2}\} \left(\sigma \sqrt{\frac{(d+T)(\log(2\delta^{-1}N(d+T)))}{n}} \right) \\
 \sqrt{\frac{c_\eta \sqrt{1 + C'/2} \sqrt{(T + d + \log(4N\delta^{-1})) (\log^{3/2}(8\delta^{-1}N(d+T)))}}{n}}. \tag{30}$$

C PROOF OF PROPOSITION 2

C.1 Concentration bounds on the arms covariance.

For every $t \in [T]$, we recall that the empirical covariance matrix for task t as

$$\widehat{\Sigma}_{t,n} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_{t,i} \mathbf{x}_{t,i}^\top \quad (31)$$

and the corresponding adapted covariance matrix (w.r.t. the filtration $(\overline{\mathcal{F}}_n)_{n \geq 0}$) as

$$\Sigma_{t,n} = \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\mathbf{x}_{t,i} \mathbf{x}_{t,i}^\top \middle| \overline{\mathcal{F}}_{i-1} \right]. \quad (32)$$

Here *adapted* means that the covariance matrix $\Sigma_{t,n}$ is fully known at time n since it depends only on the past history $\overline{\mathcal{F}}_{n-1}$.

Moreover, we use the notation $\overline{\Sigma}, \widehat{\Sigma}_n, \overline{\Sigma}_n \in \mathbb{R}^{dT \times dT}$ for the theoretical, the empirical and the adapted multi-task matrices, respectively. They are all block diagonal and composed by the corresponding T single task $d \times d$ matrices on the diagonal (e.g. $\overline{\Sigma}_n = \text{diag}(\overline{\Sigma}_{1,n}, \dots, \overline{\Sigma}_{T,n})$).

The following lemma gives a high probability bound on the deviation of $\widehat{\Sigma}_n$ from $\overline{\Sigma}_n$ according to the operator norm restricted to the cone \mathcal{C}_r defined in (7). We define for any symmetric $dT \times dT$ matrix A

$$\|A\|_{\text{op}, \mathcal{C}_r} = \max_{\Delta \in \mathcal{C}_r : \|\Delta\|_F = 1} \{ \langle A \text{Vec}(\Delta), \text{Vec}(\Delta) \rangle \}.$$

Lemma 2 (Concentration). *Let Assumption 1 be satisfied. For any $n \geq 1$ and $\delta \in (0, 1)$, with probability at least $1 - \delta$, for any $n \in [N]$*

$$\left\| \widehat{\Sigma}_n - \overline{\Sigma}_n \right\|_{\text{op}, \mathcal{C}_r} \leq C \max_{1 \leq k \leq K} \left\{ \|\Sigma_k\|_{\text{op}} \right\} r \left(\frac{(1 + \log(4NT\delta^{-1})) \log^2(d)}{\sqrt{n}} + \frac{(1 + \log(4NT\delta^{-1}))^2 \log^2(d)}{n} \right). \quad (33)$$

for some constant $C = C(C_{\mathbf{z}}) > 0$.

Proof. For any $n \in [N]$, we define the $dT \times dT$ block-diagonal matrix

$$\mathbf{M}_n := \text{diag}(\mathbf{M}_{1,n}, \dots, \mathbf{M}_{T,n}),$$

where the diagonal $d \times d$ diagonal block matrices are defined as follows:

$$\mathbf{M}_{t,n} := n \left(\widehat{\Sigma}_{t,n} - \Sigma_{t,n} \right) = \sum_{s=1}^n \mathbf{x}_{t,s} \mathbf{x}_{t,s}^\top - \mathbb{E} \left[\mathbf{x}_{t,s} \mathbf{x}_{t,s}^\top \middle| \overline{\mathcal{F}}_{s-1} \right]. \quad (34)$$

We also set $\mathbf{M}_0 = \mathbf{0}_{dT \times dT}$ a.s. By construction $(\mathbf{M}_n)_{n \geq 0}$ is a $\{\overline{\mathcal{F}}_n\}_{n \geq 0}$ -martingale satisfying

$$\|\mathbf{M}_n\|_{\text{op}, \mathcal{C}_r} = \max_{\Delta \in \mathcal{C}_r : \|\Delta\|_F = 1} \{ \langle \mathbf{M}_n \text{Vec}(\Delta), \text{Vec}(\Delta) \rangle \} = \max_{\Delta \in \mathcal{C}_r : \|\Delta\|_F = 1} \left\{ \sum_{t=1}^T \langle \mathbf{M}_{t,n} \Delta_t, \Delta_t \rangle \right\}. \quad (35)$$

We derive several geometric properties of matrices in the cone \mathcal{C}_r . We then exploit these properties via an improved union bound argument to derive a sharper uniform deviation bound of the martingale $(\mathbf{M}_n)_{n \geq 0}$ on the cone \mathcal{C}_r as compared to the standard operator norm bound on the whole space of $d \times T$ matrices.

Uniform deviation bound on $\langle \mathbf{M}_{t,n} \mathbf{u}, \mathbf{v} \rangle$ for vectors \mathbf{u}, \mathbf{v} in lower-dimensional balls. We consider two linear subspaces \mathbf{U}, \mathbf{V} of \mathbb{R}^d of respective dimension $d_{\mathbf{U}}$ and $d_{\mathbf{V}}$. We then introduce the ball $B_{\mathbf{U}}(0, r_{\mathbf{U}}) \in \mathbf{U}$ centered at 0 of radius $r_{\mathbf{U}}$. We define similarly the ball $B_{\mathbf{V}}(0, r_{\mathbf{V}}) \in \mathbf{V}$. Fix $\mathbf{u} \in B_{\mathbf{U}}(0, r_{\mathbf{U}})$ and $\mathbf{v} \in B_{\mathbf{V}}(0, r_{\mathbf{V}})$. We consider now the $\{\overline{\mathcal{F}}_n\}_{n \geq 0}$ -martingale $(\langle \mathbf{M}_{t,n} \mathbf{u}, \mathbf{v} \rangle)_{n \geq 0}$. Note that

$$\max_{\mathbf{u} \in B_{\mathbf{U}}(0, r_{\mathbf{U}}), \mathbf{v} \in B_{\mathbf{V}}(0, r_{\mathbf{V}})} \langle \mathbf{M}_{t,n} \mathbf{u}, \mathbf{v} \rangle \leq r_{\mathbf{U}} r_{\mathbf{V}} \max_{\mathbf{u} \in B_{\mathbf{U}}(0,1), \mathbf{v} \in B_{\mathbf{V}}(0,1)} \langle \mathbf{M}_{t,n} \mathbf{u}, \mathbf{v} \rangle \leq r_{\mathbf{U}} r_{\mathbf{V}} \|P_{\mathbf{U}} \mathbf{M}_{t,n} P_{\mathbf{V}}\|_{\text{op}}. \quad (36)$$

Using Lemma 4 with δ replaced by $\delta/(2T)$. gives with probability at least $1 - \delta$, for any $n \in [N]$ and any $t \in [T]$

$$\begin{aligned} & \max_{\mathbf{u} \in B_{\mathbf{U}}(0, r_{\mathbf{U}}), \mathbf{v} \in B_{\mathbf{V}}(0, r_{\mathbf{V}})} \left\langle \frac{\mathbf{M}_{t,n}}{n} \mathbf{u}, \mathbf{v} \right\rangle \\ & \leq Cr_{\mathbf{U}} r_{\mathbf{V}} \max_{k \in [K]} \{\|\Sigma_k\|_{\text{op}}\} \left(\sqrt{\frac{(d_{\mathbf{U}} \vee d_{\mathbf{V}} + \log(4NT\delta^{-1})) \log(8NT\delta^{-1}d)}{n}} \right. \\ & \quad \left. + \frac{\sqrt{(d_{\mathbf{U}} + \log(4NT\delta^{-1})) (d_{\mathbf{V}} + \log(4NT\delta^{-1}d)) \log(4NT\delta^{-1})}}{n} \right). \end{aligned} \quad (37)$$

Geometric properties of \mathcal{C}_r . Set $\bar{r} = \text{rank}(\Delta)$. Note that $\bar{r} \leq d \wedge T$. Taking the SVD of Δ , we have

$$\Delta = \sum_{j=1}^{\bar{r}} \sigma_j(\Delta) \mathbf{u}_j(\Delta) \otimes \mathbf{v}_j(\Delta),$$

with singular values $\sigma_1(\Delta) \geq \sigma_2(\Delta) \geq \dots \geq \sigma_{\bar{r}}(\Delta) > 0$ and orthonormal families $\{\mathbf{u}_j(\Delta)\}_{j=1}^{\bar{r}} \in \mathbb{R}^d$, $\{\mathbf{v}_j(\Delta)\}_{j=1}^{\bar{r}} \in \mathbb{R}^T$. For the sake of brevity, we set $\sigma_j(\Delta) = \sigma_j$, $\mathbf{u}_j(\Delta) = \mathbf{u}_j$ and $\mathbf{v}_j(\Delta) = \mathbf{v}_j$ for any $j \in [\bar{r}]$. From the previous displays, we immediately get the following representation for the columns of Δ :

$$\Delta_t = \sum_{j=1}^{\bar{r}} \sigma_j \mathbf{v}_{j,t} \mathbf{u}_j, \quad \forall t \in T, \quad (38)$$

where \mathbf{v}_j admits components $\mathbf{v}_j = (\mathbf{v}_{j,1}, \dots, \mathbf{v}_{j,T})^\top$.

By definition of the cone \mathcal{C}_r in (7), for any $\Delta \in \mathcal{C}_r$, we have $\|\Pi(\Delta)\|_* \leq 3\|\Delta - \Pi(\Delta)\|_*$. Note that $\text{rank}(\Delta - \Pi(\Delta)) \leq 2r$ by definition of Π and the cone \mathcal{C}_r . Hence we have, for any $\Delta \in \mathcal{C}_r$ with $\|\Delta\|_F = 1$,

$$\|\Delta\|_* \leq \|\Pi(\Delta)\|_* + \|\Delta - \Pi(\Delta)\|_* \leq 4\|\Delta - \Pi(\Delta)\|_* \leq 4\sqrt{2r}\|\Delta - \Pi(\Delta)\|_F \leq 4\sqrt{2r}\|\Delta\|_F = 4\sqrt{2r}. \quad (39)$$

We deduce from (39) that $\sum_{j=1}^{\bar{r}} \sigma_j \leq 4\sqrt{2r}$ and consequently for any $\Delta \in \mathcal{C}_r$

$$\sigma_j(\Delta) \leq \frac{4\sqrt{2r}}{j}, \quad \forall j \in [\bar{r}]. \quad (40)$$

We conclude this paragraph with some elementary facts on orthonormal basis. We complete the orthonormal family $\{\mathbf{v}_j\}_{j=1}^{\bar{r}}$ into an orthonormal basis $\{\mathbf{v}_j\}_{j=1}^T$ of \mathbb{R}^T . By properties of orthonormal basis, we have $\langle \mathbf{v}_j, \mathbf{v}_{j'} \rangle = \sum_{t=1}^T \mathbf{v}_{j,t} \mathbf{v}_{j',t} = \delta_{j,j'}$ where $\delta_{j,j'} = 1$ if $j = j'$ and 0 otherwise.

Uniform bound over $\Delta \in \mathcal{C}_r$. Our goal is to control $\|\mathbf{M}_n\|_{\text{op}, \mathcal{C}_r}$. To this end, we set $m_* = \lceil \log_2(d \wedge T) \rceil$ and, for any $m \in [m_*]$, define $J_m = \{j \in [m_*] : 2^{m-1} \leq j < 2^m\}$. Next we propose the following decomposition of Δ :

$$\Delta_t = \sum_{m=1}^{m_*} \Delta_{t,m}, \quad \text{where} \quad \Delta_{t,m} = \sum_{j \in J_m} \sigma_j \mathbf{v}_{j,t} \mathbf{u}_j, \quad \forall t \in [T].$$

By construction and (40), for any $m \in [m_*]$, all the vectors $\Delta_{t,m}$, $t \in [T]$ live in the same subspace of dimension at most 2^{m-1} and

$$\|\Delta_{t,m}\|^2 \leq \sigma_{2^{m-1}}^2 \sum_{j \in J_m} \mathbf{v}_{j,t}^2 \leq \sigma_{2^{m-1}}^2 \leq \frac{32r}{2^{2(m-1)}},$$

and

$$\sum_{t=1}^T \|\Delta_{t,m}\|^2 \leq \sigma_{2^{m-1}}^2 \sum_{j \in J_m} \sum_{t=1}^T \mathbf{v}_{j,t}^2 \leq \sigma_{2^{m-1}}^2 |J_m| \leq \frac{32r}{2^{(m-1)}}. \quad (41)$$

Combining the last two displays with (38), we get that

$$\sum_{t=1}^T \langle \mathbf{M}_{t,n} \Delta_t, \Delta_t \rangle = \sum_{t=1}^T \sum_{m,m'=1}^{m_*} \langle \mathbf{M}_{t,n} \Delta_{t,m}, \Delta_{t,m'} \rangle. \quad (42)$$

An union bound combining the last two displays with (35) and (37) gives for any $\delta \in (0, 1)$ with probability at least $1 - \delta$, for any $n \in [N]$

$$\frac{1}{n} \|\mathbf{M}_n\|_{\text{op}, \mathcal{C}_r} \leq C \max_{1 \leq k \leq K} \left\{ \|\Sigma_k\|_{\text{op}} \right\} \sum_{m,m'=1}^{m_*} \sum_{t=1}^T \|\Delta_{t,m}\| \|\Delta_{t,m'}\| \epsilon_{m,m'},$$

where

$$\begin{aligned} \epsilon_{m,m'} = & \sqrt{\frac{(2^{m \vee m'} - 1 + \log(4NT\delta^{-1})) \log(4NT\delta^{-1}d)}{n}} \\ & + \frac{\sqrt{(2^{m-1} + \log(4NT\delta^{-1})) (2^{m'-1} + \log(4NT\delta^{-1}d)) \log(4NT\delta^{-1})}}{n}. \end{aligned}$$

The Cauchy-Schwartz inequality and (41) give

$$\begin{aligned} \|\mathbf{M}_n\|_{\text{op}, \mathcal{C}_r} & \leq C \max_{1 \leq k \leq K} \left\{ \|\Sigma_k\|_{\text{op}} \right\} \sum_{m,m'=1}^{m_*} \left(\sum_{t=1}^T \|\Delta_{t,m}\|^2 \right)^{1/2} \left(\sum_{t=1}^T \|\Delta_{t,m'}\|^2 \right)^{1/2} \epsilon_{m,m'} \\ & \leq 32C \max_{1 \leq k \leq K} \left\{ \|\Sigma_k\|_{\text{op}} \right\} r \sum_{m,m'=1}^{m_*} \frac{1}{2^{\frac{m-1}{2}}} \frac{1}{2^{\frac{m'-1}{2}}} \epsilon_{m,m'}. \end{aligned}$$

By definition of $\epsilon_{m,m'}$, we have for any $m, m' \in [m_*]$

$$\frac{1}{2^{\frac{m-1}{2}}} \frac{1}{2^{\frac{m'-1}{2}}} \epsilon_{m,m'} \lesssim \frac{1 + \log(4NT\delta^{-1})}{\sqrt{n}} + \frac{(1 + \log(4NT\delta^{-1}))^2}{n}.$$

Combining the last two displays with the fact that $m_* \lesssim \log(d)$, we get for any $\delta \in (0, 1)$ with probability at least $1 - \delta$, for any $n \in [N]$

$$\frac{1}{n} \|\mathbf{M}_n\|_{\text{op}, \mathcal{C}_r} \leq C \max_{1 \leq k \leq K} \left\{ \|\Sigma_k\|_{\text{op}} \right\} r \left(\frac{(1 + \log(4NT\delta^{-1})) \log^2(d)}{\sqrt{n}} + \frac{(1 + \log(4NT\delta^{-1}))^2 \log^2(d)}{n} \right).$$

□

C.2 RSC condition and Matrix perturbation

The next lemma guarantees that if the $\|\cdot\|_{\text{op}, \mathcal{C}_r}$ metric between two matrices is small enough and if one of these two matrices satisfies the RSC condition, then the other matrix also satisfied the RSC condition. Specifically, the next lemma links the RSC constant $\kappa(\overline{\Sigma})$ associated to the multi-task theoretical matrix, to the adapted one $\kappa(\overline{\Sigma}_n)$.

Lemma 3 (RSC condition and Random Matrices). *Let Σ_0 and Σ_1 be two $dT \times dT$ block-diagonal matrices (with blocks matrices of size $d \times d$). Suppose that the RSC condition is met for the covariance matrix Σ_0 with constant $\kappa(\Sigma_0)$ and that $\|\Sigma_0 - \Sigma_1\|_{\text{op}, \mathcal{C}_r} \leq \tilde{\lambda}$, where*

$$\tilde{\lambda} \leq \kappa(\Sigma_0).$$

Then, the RSC condition is also met for the covariance matrix Σ_1 with constant $\kappa(\Sigma_1) \geq \kappa(\Sigma_0)/2$. Moreover, for all $\Delta \in \mathbb{R}^{d \times T}$ such that $\Delta \in \mathcal{C}(r)$ (see equation 8), we have

$$\|\text{Vec}(\Delta)\|_{\Sigma_0}^2 \leq 2 \|\text{Vec}(\Delta)\|_{\Sigma_1}^2 \leq 3 \|\text{Vec}(\Delta)\|_{\Sigma_0}^2. \quad (43)$$

The proof combines Definition 1 and the $\|\cdot\|_{\text{op}, \mathcal{C}_r}$ metric closeness $\|\Sigma_0 - \Sigma_1\|_{\text{op}, \mathcal{C}_r} \leq \tilde{\lambda}$.

Proof. For any $\Delta \in \mathcal{C}_r$, we have

$$\left| \|\text{Vec}(\Delta)\|_{\Sigma_0}^2 - \|\text{Vec}(\Delta)\|_{\Sigma_1}^2 \right| \leq \|\Sigma_0 - \Sigma_1\|_{\text{op}, \mathcal{C}_r} \|\text{Vec}(\Delta)\|_2^2 \leq \tilde{\lambda} \|\text{Vec}(\Delta)\|_2^2.$$

According to Definition 1, for any $\Delta \in \mathcal{C}_r$, the RSC condition gives

$$\|\Delta\|_{\text{F}}^2 = \|\text{Vec}(\Delta)\|_2^2 \leq \frac{\|\text{Vec}(\Delta)\|_{\Sigma_0}^2}{2\kappa(\Sigma_0)}.$$

Combining these two last results we have

$$\left| \|\text{Vec}(\Delta)\|_{\Sigma_0}^2 - \|\text{Vec}(\Delta)\|_{\Sigma_1}^2 \right| \leq \tilde{\lambda} \frac{\|\text{Vec}(\Delta)\|_{\Sigma_0}^2}{2\kappa(\Sigma_0)}.$$

Equivalently, the following result holds

$$\left| \frac{\|\text{Vec}(\Delta)\|_{\Sigma_1}^2}{\|\text{Vec}(\Delta)\|_{\Sigma_0}^2} - 1 \right| \leq \frac{\tilde{\lambda}}{2\kappa(\Sigma_0)}.$$

The statement follows directly by the selected $\tilde{\lambda}$. □

C.3 Combining everything

We can now prove Proposition 2.

Proof. In the RSC condition (see Definition 1) the considered covariance matrix $\bar{\Sigma} \in \mathbb{R}^{dT \times dT}$ is block-diagonal consisting of T blocks, one for each different task. Let us consider now the single-task adapted matrix

$$\Sigma_{t,n} = \frac{1}{n} \sum_{s=1}^n \mathbb{E} \left[\mathbf{x}_{t,s} \mathbf{x}_{t,s}^\top \middle| \bar{\mathcal{F}}_{s-1} \right]$$

and the multi-task one $\bar{\Sigma}_n = \text{diag}(\Sigma_{1,n}, \dots, \Sigma_{T,n}) \in \mathbb{R}^{dT \times dT}$. Relying on (Oh et al., 2021, Lemma 10), under Assumption 2 we have

$$\bar{\Sigma}_n = \frac{1}{n} \sum_{s=1}^n \mathbb{E} \left[\bar{\mathbf{x}}_s \bar{\mathbf{x}}_s^\top \middle| \bar{\mathcal{F}}_{s-1} \right] \succeq (2\nu\omega_{\mathcal{X}})^{-1} \bar{\Sigma}, \quad (44)$$

where $\bar{\mathbf{x}}_s = [\mathbf{x}_{1,s}^\top, \dots, \mathbf{x}_{T,s}^\top]^\top \in \mathbb{R}^{dT}$. Now, let us denote $\tilde{\mathbf{w}}_n = \arg \min_{\tilde{\mathbf{w}} \in \mathcal{C}(r)} \frac{\text{Vec}(\tilde{\mathbf{w}})^\top \bar{\Sigma}_n \text{Vec}(\tilde{\mathbf{w}})}{2\|\text{Vec}(\tilde{\mathbf{w}})\|_2^2}$.

Relying on the RSC condition and thanks to the previous display, we obtain

$$\frac{\text{Vec}(\tilde{\mathbf{w}})_n^\top \bar{\Sigma}_n \text{Vec}(\tilde{\mathbf{w}})_n}{\|\text{Vec}(\tilde{\mathbf{w}})_n\|_2^2} \geq \frac{\text{Vec}(\tilde{\mathbf{w}})_n^\top \bar{\Sigma} \text{Vec}(\tilde{\mathbf{w}})_n}{2\nu\omega_{\mathcal{X}} \|\text{Vec}(\tilde{\mathbf{w}})_n\|_2^2} \geq \frac{\kappa(\bar{\Sigma})}{2\nu\omega_{\mathcal{X}}}.$$

Hence, $\bar{\Sigma}_n$ satisfies the RSC condition with constant $\kappa(\bar{\Sigma}_n) = \kappa(\bar{\Sigma}) / (2\nu\omega_{\mathcal{X}})$.

Next we consider the operator norm deviation of the multi-task adapted matrix $\Sigma_0 = \bar{\Sigma}_n$ from the multi-task empirical matrix $\Sigma_1 = \hat{\Sigma}_n$. In view of (33), the quantity $\left\| \hat{\Sigma}_n - \bar{\Sigma}_n \right\|_{\text{op}, \mathcal{C}_r}$ becomes smaller as n increases. This means that if we take n large enough, then we have whp that $\left\| \hat{\Sigma}_n - \bar{\Sigma}_n \right\|_{\text{op}, \mathcal{C}_r} \leq \kappa(\bar{\Sigma}) / (4\nu\omega_{\mathcal{X}}) = \kappa(\bar{\Sigma}_n) / 2$. We formalize this argument below.

For a fixed $\delta \in (0, 1)$, recall the definition of $N_0(\delta)$ in (12). In view of Lemma 2, if $n \geq N_0(\delta)$ then with probability at least $1 - \delta$,

$$\left\| \hat{\Sigma}_n - \bar{\Sigma}_n \right\|_{\text{op}, \mathcal{C}_r} \leq \frac{\kappa(\bar{\Sigma})}{4\nu\omega_{\mathcal{X}}} \leq \frac{\kappa(\bar{\Sigma}_n)}{2}.$$

Then Lemma 3 yields the result. □

C.4 Technical Lemma

Consider linear subspaces \mathbf{U}, \mathbf{V} of \mathbb{R}^d of dimension $d_{\mathbf{U}}$ and $d_{\mathbf{V}}$ respectively. Denote by $P_{\mathbf{U}}$ and $P_{\mathbf{V}}$ the orthogonal projections onto \mathbf{U} and \mathbf{V} respectively. We have for any $n \in [N]$

$$\max_{\mathbf{u} \in B_{\mathbf{U}}(0,1), \mathbf{v} \in B_{\mathbf{V}}(0,1)} \langle \mathbf{M}_{t,n} \mathbf{u}, \mathbf{v} \rangle \leq \|P_{\mathbf{U}} \mathbf{M}_{t,n} P_{\mathbf{V}}\|_{\text{op}}.$$

Lemma 4. *Let Assumption 1 be satisfied. For any $\delta \in (0, 1)$ and any $t \in [T]$, with probability at least $1 - \delta$,*

$$\max_{n \in [N]} \left\{ \frac{1}{n} \|P_{\mathbf{U}} \mathbf{M}_{t,n} P_{\mathbf{V}}\|_{\text{op}} \right\} \leq C \max_{k \in [K]} \{ \|\Sigma_k\|_{\text{op}} \} \left(\sqrt{\frac{(d_{\mathbf{U}} \vee d_{\mathbf{V}} + \log(4N\delta^{-1})) \log(4N\delta^{-1}d)}{n}} \right. \\ \left. + \frac{\sqrt{(d_{\mathbf{U}} + \log(4N\delta^{-1})) (d_{\mathbf{V}} + \log(4N\delta^{-1}d)) \log(4N\delta^{-1})}}{n} \right), \quad (45)$$

where $C = C(C_{\mathbf{z}}) > 0$ may depend only on $C_{\mathbf{z}}$.

Proof. We have for any $n \geq 1$

$$\|P_{\mathbf{U}}(\mathbf{M}_{t,n} - \mathbf{M}_{t,n-1})P_{\mathbf{V}}\|_{\text{op}} = \left\| P_{\mathbf{U}}(\mathbf{x}_{t,n})P_{\mathbf{V}}(\mathbf{x}_{t,n})^{\top} - \mathbb{E} \left[P_{\mathbf{U}}(\mathbf{x}_{t,n})P_{\mathbf{V}}(\mathbf{x}_{t,n})^{\top} \middle| \overline{\mathcal{F}}_{n-1} \right] \right\|_{\text{op}} \\ \leq \|P_{\mathbf{U}}(\mathbf{x}_{t,n})\| \|P_{\mathbf{V}}(\mathbf{x}_{t,n})\| + \max_{k \in [K]} \{ \|\Sigma_k\|_{\text{op}} \}. \quad (46)$$

In view of Assumption 1, we have $\|P_{\mathbf{U}}(\mathbf{x}_{t,n})\|_2^2 = \mathbf{z}_{t,n}^{\top} P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}} \mathbf{z}_{t,n}$ for some $k \in [K]$. We apply now Hanson-Wright's inequality conditionally on $\overline{\mathcal{F}}_{n-1}$ to get for any $x > 0$

$$\mathbb{P} \left(\mathbf{z}_{t,n}^{\top} P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}} \mathbf{z}_{t,n} \leq \mathbb{E} [\mathbf{z}_{t,n}^{\top} P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}} \mathbf{z}_{t,n} | \overline{\mathcal{F}}_{n-1}] + C C_{\mathbf{z}}^2 \left(\sqrt{\|P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}}\|_{\mathbb{F}}^2 x} + \|P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}}\|_{\text{op}} x \right) \middle| \overline{\mathcal{F}}_{n-1} \right) \geq 1 - e^{-x}.$$

Note that $\mathbb{E} [\mathbf{z}_{t,n}^{\top} P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}} \mathbf{z}_{t,n} | \overline{\mathcal{F}}_{n-1}] = \text{trace}(P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}}) \leq \max_{k \in [K]} \{ \|\Sigma_k\|_{\text{op}} \} d_{\mathbf{U}}$ since \mathbf{z}_k is a zero mean isotropic random vector. We define the following event in $\overline{\mathcal{F}}_n$

$$\Omega_n(\mathbf{U}) := \left\{ \|P_{\mathbf{U}}(\mathbf{x}_{t,n})\|^2 \leq \max_{k \in [K]} \{ \|\Sigma_k\|_{\text{op}} \} \left(d_{\mathbf{U}} + C C_{\mathbf{z}}^2 \left(\sqrt{d_{\mathbf{U}} \log(4N\delta^{-1})} + \log(4N\delta^{-1}) \right) \right) \right\}, \quad (47)$$

and

$$\Omega_n(\mathbf{V}) := \left\{ \|P_{\mathbf{V}}(\mathbf{x}_{t,n})\|^2 \leq \max_{k \in [K]} \{ \|\Sigma_k\|_{\text{op}} \} \left(d_{\mathbf{V}} + C C_{\mathbf{z}}^2 \left(\sqrt{d_{\mathbf{V}} \log(4N\delta^{-1})} + \log(4N\delta^{-1}) \right) \right) \right\}. \quad (48)$$

Combining the last three displays and taking $x = \log(\delta^{-1})$ immediately implies for any $n \in [N]$

$$\mathbb{P}(\Omega_n(\mathbf{U}) | \overline{\mathcal{F}}_{n-1}) \geq 1 - \frac{\delta}{4N}, \quad \mathbb{P}(\Omega_n(\mathbf{V}) | \overline{\mathcal{F}}_{n-1}) \geq 1 - \frac{\delta}{4N} \quad (49)$$

We define, for any $n \in [N]$

$$\overline{\Omega}_n = \bigcap_{l=0}^n \Omega_l(\mathbf{U}) \cap \Omega_l(\mathbf{V}) \in \overline{\mathcal{F}}_n, \quad (50)$$

with $\Omega_0(\mathbf{U}) = \Omega_0(\mathbf{V})$ being the whole sample space. The Bayes rule combined with (49) gives

$$\mathbb{P}(\overline{\Omega}_N) = \prod_{n=1}^N \mathbb{P} \left(\Omega_n(\mathbf{U}) \cap \Omega_n(\mathbf{V}) \middle| \bigcap_{k=0}^{n-1} \Omega_k(\mathbf{U}) \cap \Omega_k(\mathbf{V}) \right) \geq \left(1 - \frac{\delta}{2N} \right)^N. \quad (51)$$

Bernoulli's inequality ($(1+x)^n \geq 1+nx$ for any $x > -1$ and integer $n \geq 1$) gives

$$\mathbb{P}(\overline{\Omega}_N) = \prod_{n=1}^N \mathbb{P} \left(\Omega_n(\mathbf{U}) \cap \Omega_n(\mathbf{V}) \middle| \bigcap_{k=0}^{n-1} \Omega_k(\mathbf{U}) \cap \Omega_k(\mathbf{V}) \right) \geq 1 - \frac{\delta}{2}. \quad (52)$$

Define the stopping times $\tau(\mathbf{U})$ and $\tau(\mathbf{V})$ as follows

$$\tau(\mathbf{U}) = \inf \left\{ n \geq 0 : \|P_{\mathbf{U}}(\mathbf{x}_{t,n})\|^2 \geq \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}} \right\} \left(d_{\mathbf{U}} + C C_{\mathbf{z}}^2 \left(\sqrt{d_{\mathbf{U}} \log(4N\delta^{-1})} + \log(4N\delta^{-1}) \right) \right) \right\}, \quad (53)$$

and

$$\tau(\mathbf{V}) = \inf \left\{ n \geq 0 : \|P_{\mathbf{V}}(\mathbf{x}_{t,n})\|^2 \geq \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}} \right\} \left(d_{\mathbf{V}} + C C_{\mathbf{z}}^2 \left(\sqrt{d_{\mathbf{V}} \log(4N\delta^{-1})} + \log(4N\delta^{-1}) \right) \right) \right\}, \quad (54)$$

Again by definition of the Trace-Norm bandit, $\tau(\mathbf{U})$ and $\tau(\mathbf{V})$ are both stopping times relative to the filtration $\{\bar{\mathcal{F}}_n\}_{n \geq 0}$, so is $\tau = \tau(\mathbf{U}) \wedge \tau(\mathbf{V})$. Hence the stopped process $\{\bar{\mathbf{M}}_n^\tau\}_{n \geq 0}$ defined as follows

$$\bar{\mathbf{M}}_n^\tau := P_{\mathbf{U}} \mathbf{M}_{t,n \wedge \tau} P_{\mathbf{V}}, \quad \forall n \geq 0,$$

is also a martingale adapted to the filtration $\{\bar{\mathcal{F}}_n\}_{n \geq 0}$. Furthermore $\{\bar{\mathbf{M}}_n^\tau\}_{n \geq 0}$ is a bounded martingale. Hence we can apply the Freedman inequality to it.

By definition of τ and (52), we have $\{\tau > N\} = \bar{\Omega}_N$. Hence for any $x > 0$

$$\begin{aligned} \mathbb{P} \left(\|P_{\mathbf{U}} \mathbf{M}_{t,n} P_{\mathbf{V}}\|_{\text{op}} \geq x \right) &= \mathbb{P} \left(\{P_{\mathbf{U}} \mathbf{M}_{t,n} P_{\mathbf{V}} \geq x\} \cap \{\tau > N\} \right) + \mathbb{P} \left(\{P_{\mathbf{U}} \mathbf{M}_{t,n} P_{\mathbf{V}}\|_{\text{op}} \geq x\} \cap \{\tau \leq N\} \right) \\ &\leq \mathbb{P} \left(\|\bar{\mathbf{M}}_n^\tau\|_{\text{op}} \geq x \right) + \mathbb{P}(\{\tau \leq N\}) \leq \mathbb{P} \left(\|\bar{\mathbf{M}}_n^\tau\|_{\text{op}} \geq x \right) + \delta/2, \end{aligned} \quad (55)$$

where we have used (52) in the last line.

Deviation bound for $\|\bar{\mathbf{M}}_n^\tau\|_{\text{op}}$. We will use again (Tropp, 2011, Corollary 1.3).

In view of (46), (47) and (48), we have

$$\|\bar{\mathbf{M}}_n^\tau - \bar{\mathbf{M}}_{n-1}^\tau\|_{\text{op}} \leq \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}} \right\} \left(1 + (1 + C C_{\mathbf{z}}^2/2) \sqrt{(d_{\mathbf{U}} + \log(4N\delta^{-1})) (d_{\mathbf{V}} + \log(4N\delta^{-1}))} \right). \quad (56)$$

Next, we set $Y_{t,n} = \bar{\mathbf{M}}_n^\tau - \bar{\mathbf{M}}_{n-1}^\tau$. Elementary computations give

$$\mathbf{W}_{\text{col}} = \sum_{s=1}^n \mathbb{E} [Y_{t,n} Y_{t,n}^\top | \bar{\mathcal{F}}_{s-1}] \leq (1 + C C_{\mathbf{z}}^2/2)^2 n (d_{\mathbf{V}} + \log(4N\delta^{-1})) \max_{k \in [K]} \{ \|P_{\mathbf{V}} \Sigma_k P_{\mathbf{V}}\|_{\text{op}} \} \max_{k \in [K]} \{ \|P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}}\|_{\text{op}} \},$$

and

$$\mathbf{W}_{\text{row}} = \sum_{s=1}^n \mathbb{E} [Y_{t,n}^\top Y_{t,n} | \bar{\mathcal{F}}_{s-1}] \leq (1 + C C_{\mathbf{z}}^2/2)^2 n (d_{\mathbf{U}} + \log(4N\delta^{-1})) \max_{k \in [K]} \{ \|P_{\mathbf{V}} \Sigma_k P_{\mathbf{V}}\|_{\text{op}} \} \max_{k \in [K]} \{ \|P_{\mathbf{U}} \Sigma_k P_{\mathbf{U}}\|_{\text{op}} \}.$$

Hence

$$\mathbf{W}_{\text{row}} \vee \mathbf{W}_{\text{col}} \leq (1 + C C_{\mathbf{z}}^2/2)^2 n (d_{\mathbf{U}} \vee d_{\mathbf{V}} + \log(4N\delta^{-1})) \max_{k \in [K]} \{ \|\Sigma_k\|_{\text{op}}^2 \}. \quad (57)$$

Applying (Tropp, 2011, Corollary 1.3), we get for any $\delta \in (0, 1)$ with probability at least $1 - \delta/(2N)$ that

$$\begin{aligned} \|\bar{\mathbf{M}}_n^\tau\|_{\text{op}} &\leq (1 + C C_{\mathbf{z}}^2/2) \max_{k \in [K]} \{ \|\Sigma_k\|_{\text{op}} \} \left(\sqrt{n (d_{\mathbf{U}} \vee d_{\mathbf{V}} + \log(4N\delta^{-1})) \log(4N\delta^{-1}d)} \right. \\ &\quad \left. + 2 \sqrt{(d_{\mathbf{U}} + \log(4N\delta^{-1})) (d_{\mathbf{V}} + \log(4N\delta^{-1})) \log(4N\delta^{-1}d)} \right). \end{aligned} \quad (58)$$

Dividing by n and a trivial union bound gives the result. \square

D PROOF OF THEOREM 1

We recall that $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_T]$ is of rank smaller than r . Hence \mathbf{W} admits the following singular value decomposition

$$\mathbf{W} = \sum_{j=1}^r \sigma_j(\mathbf{W}) u_j(\mathbf{W}) \otimes v_j(\mathbf{W}),$$

with singular values $\sigma_1(\mathbf{W}) \geq \dots \geq \sigma_r(\mathbf{W}) > 0$ and corresponding left and right orthonormal singular vectors $u_1(\mathbf{W}), \dots, u_r(\mathbf{W})$ of \mathbb{R}^d and $v_1(\mathbf{W}), \dots, v_r(\mathbf{W})$ of \mathbb{R}^T . We denote by P_U the orthogonal projection onto $U = \text{l.s.}(u_1(\mathbf{W}), \dots, u_r(\mathbf{W}))$. By definition of the SVD, we obviously have $P_U(\mathbf{w}_t) = \mathbf{w}_t$ for any $1 \leq t \leq T$.

We consider now the instantaneous multi-task regret

$$\bar{R}_n = \sum_{t=1}^T R_{t,n} = \sum_{t=1}^T \langle \mathbf{x}_{t,n}^*, \mathbf{w}_t \rangle - \langle \mathbf{x}_{t,n}, \mathbf{w}_t \rangle \leq 2 \sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{D}_{t,n}} |\langle \mathbf{x}, \mathbf{w}_t \rangle|,$$

where $\mathbf{x}_{t,n}^* = \arg \max_{\mathbf{x} \in \mathcal{D}_{t,n}} \langle \mathbf{x}, \mathbf{w}_t \rangle$.

Assumption 1 guarantees that $\frac{\langle \mathbf{x}, \mathbf{w}_t \rangle}{\|\mathbf{w}_t\|}$ is $C_{\mathbf{x}}$ -subgaussian and consequently, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$

$$|\langle \mathbf{x}, \mathbf{w}_t \rangle| \leq C_{\mathbf{x}} \|\mathbf{w}_t\| \sqrt{\log(e\delta^{-1})} \leq C_{\mathbf{z}} \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} L \sqrt{\log(e\delta^{-1})}.$$

An union bound over $t \in [T]$, $n \in [N]$, $\mathbf{x} \in \mathcal{D}_{t,n}$ gives with probability at least $1 - \delta/2$

$$\max_{\mathbf{x} \in \mathcal{D}_{t,n}} |\langle \mathbf{x}, \mathbf{w}_t \rangle| \leq C_{\mathbf{z}} \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} L \sqrt{\log(2eTKN\delta^{-1})}, \quad \forall t \in [T], \forall k \in [K], \forall n \in [N].$$

Hence with the same probability, $\forall n \in [N]$

$$\bar{R}_n \leq C_{\mathbf{z}} T \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} L \sqrt{\log(2eTKN\delta^{-1})}$$

We will now split the regret into two phases. During the first N_0 interactions the collected data will not meet the RSC condition. Hence, the incurred regret will be of order $TN_0 \sqrt{\log(2eTKN\delta^{-1})}$. Conversely, from the $N_0 + 1$ -th round up to the N -th one, the RSC condition will be met allowing the result of Lemma 1.

We have

$$\begin{aligned} \sum_{n=\lceil N_0 \rceil}^N \bar{R}_n &= \sum_{n=\lceil N_0 \rceil}^N \sum_{t=1}^T \langle \mathbf{x}_{t,n}^*, \mathbf{w}_t - \hat{\mathbf{w}}_{t,n} \rangle + \sum_{n=\lceil N_0 \rceil}^N \sum_{t=1}^T \langle \mathbf{x}_{t,n}^* - \mathbf{x}_{t,n}, \hat{\mathbf{w}}_{t,n} \rangle + \sum_{n=\lceil N_0 \rceil}^N \sum_{t=1}^T \langle \mathbf{x}_{t,n}, \hat{\mathbf{w}}_{t,n} - \mathbf{w}_t \rangle \\ &\leq \sum_{n=\lceil N_0 \rceil}^N \sum_{t=1}^T \langle \mathbf{x}_{t,n}^*, \mathbf{w}_t - \hat{\mathbf{w}}_{t,n} \rangle + \sum_{n=\lceil N_0 \rceil}^N \sum_{t=1}^T \langle \mathbf{x}_{t,n}, \hat{\mathbf{w}}_{t,n} - \mathbf{w}_t \rangle = \sum_{n=\lceil N_0 \rceil}^N \sum_{t=1}^T \langle \mathbf{x}_{t,n}^* - \mathbf{x}_{t,n}, \mathbf{w}_t - \hat{\mathbf{w}}_{t,n} \rangle \end{aligned} \quad (59)$$

where we have used that for any $1 \leq n \leq N$, $1 \leq t \leq T$, $\langle \mathbf{x}_{t,n}^* - \mathbf{x}_{t,n}, \hat{\mathbf{w}}_{t,n} \rangle \leq 0$ a.s. by definition of our policy.

Next, we have

$$\left| \sum_{t=1}^T \langle \mathbf{x}_{t,n}^* - \mathbf{x}_{t,n}, \mathbf{w}_t - \hat{\mathbf{w}}_{t,n} \rangle \right| \leq 2 \sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{D}_{t,n}} |\langle \mathbf{x}, \mathbf{w}_t - \hat{\mathbf{w}}_{t,n} \rangle|. \quad (60)$$

In view of (4), we note that $\hat{\mathbf{W}}_n \in \bar{\mathcal{F}}_{n-1}$. Thus, conditionally on $\bar{\mathcal{F}}_{n-1}$, Assumption 1 guarantees that $\frac{\langle \mathbf{x}, \mathbf{w}_t - \hat{\mathbf{w}}_{t,n} \rangle}{\|\mathbf{w}_t - \hat{\mathbf{w}}_{t,n}\|}$ is $C_{\mathbf{x}}$ -subgaussian and consequently, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$

$$|\langle \mathbf{x}, \mathbf{w}_t - \hat{\mathbf{w}}_{t,n} \rangle| \leq C_{\mathbf{x}} \|\mathbf{w}_t - \hat{\mathbf{w}}_{t,n}\| \sqrt{\log(e\delta^{-1})} \leq C_{\mathbf{z}} \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} \|\mathbf{w}_t - \hat{\mathbf{w}}_{t,n}\| \sqrt{\log(e\delta^{-1})}.$$

An union bound over t, n, k gives with probability at least $1 - \delta$

$$\max_{t \in [T], n \in [N]} \max_{\mathbf{x} \in \mathcal{D}_{t,n}} |\langle \mathbf{x}, \mathbf{w}_t - \widehat{\mathbf{w}}_{t,n} \rangle| \leq C_{\mathbf{z}} \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} \|\mathbf{w}_t - \widehat{\mathbf{w}}_{t,n}\| \sqrt{\log(eTKN\delta^{-1})}. \quad (61)$$

Combining the last two displays, we get with probability at least $1 - \delta$, for any $t \in [T], n \in [N]$,

$$\begin{aligned} \left| \sum_{t=1}^T \langle \mathbf{x}_{t,n}^* - \mathbf{x}_{t,n}, \mathbf{w}_t - \widehat{\mathbf{w}}_{t,n} \rangle \right| &\leq 2C_{\mathbf{z}} \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} \left(\sum_{t=1}^T \|\mathbf{w}_t - \widehat{\mathbf{w}}_{t,n}\| \right) \sqrt{\log(eTKN\delta^{-1})} \\ &\leq 2C_{\mathbf{z}} \max_{k \in [K]} \left\{ \|\Sigma_k\|_{\text{op}}^{1/2} \right\} \sqrt{T} \|\mathbf{W} - \widehat{\mathbf{W}}_n\|_F \sqrt{\log(eTKN\delta^{-1})}. \end{aligned}$$

Recall that Lemmas 1 and 2 combined with Propositions 1 and 2 gives, for any $\lceil N_0 \rceil \leq n \leq N$ and any $\delta \in (0, 1)$, with probability at least $1 - \frac{\delta}{2N}$,

$$\left\| \widehat{\mathbf{W}}_n - \mathbf{W} \right\|_F \lesssim_{C_{\eta}, \sigma, C_{\mathbf{x}}, \kappa(\overline{\Sigma}), \max_{1 \leq k \leq K} \{\|\Sigma_k\|_{\text{op}}^{1/2}\}} \sqrt{\frac{r}{n}} \left(\sqrt{T + d + \log\left(\frac{4N(d+T)}{\delta}\right)} \log^{3/2}\left(\frac{8N(d+T)}{\delta}\right) \right). \quad (62)$$

Combining the last two displays, we obtain with probability at least $1 - \delta$

$$\left| \sum_{t=1}^T \langle \mathbf{x}_{t,n}^* - \mathbf{x}_{t,n}, \mathbf{w}_t - \widehat{\mathbf{w}}_{t,n+1} \rangle \right| \lesssim \sqrt{\frac{rT}{n}} \left(\sqrt{T + d + \log\left(\frac{4N(d+T)}{\delta}\right)} \log^{3/2}\left(\frac{8N(d+T)}{\delta}\right) \right) \sqrt{\log(2e\delta^{-1}TKN)}. \quad (63)$$

Summing over n in the previous display along with an union bound argument and (59) gives with probability at least $1 - \delta$

$$\sum_{n=\lceil N_0 \rceil}^N \overline{R}_s \leq C \sqrt{rTN} \left(\sqrt{T + d + \log\left(\frac{4N(d+T)}{\delta}\right)} \log^{3/2}\left(\frac{8N(d+T)}{\delta}\right) \right) \sqrt{\log(2e\delta^{-1}TKN)}, \quad (64)$$

where $C = C\left(\eta, \sigma, C_{\mathbf{z}}, \kappa(\overline{\Sigma}), \max_{1 \leq k \leq K} \{\|\Sigma_k\|_{\text{op}}^{1/2}\}\right)$ is a finite constant under our assumptions.

E NUMERICAL EXPERIMENTS

We perform here some additional experiments on the impact of the rank r , the dimension d and the number of tasks T on the performance of our trace norm bandit policy (Algorithm 1). We also compare its performance to the MLingreedy policy of Yang et al. (2020).

Comparison to MLingreedy (Yang et al., 2020). We compare the performance of the trace norm bandit (Alg. 1) to that of the MLingreedy policy and independent task learning (ITL). We recall that the MLingreedy policy use a rank parameter \bar{r} . This parameter is set equal to the true rank r of \mathbf{W} in Yang et al. (2020) to derive the theoretical properties of this policy.

In our experiments we will use different values for the rank parameter \bar{r} , either the true rank r or an overestimated rank $\bar{r} = \min(2r, d, T)$ or an underestimated rank $\bar{r} = \max(\lfloor r/2 \rfloor, 1)$. At each epoch of the the MLingreedy policy, we used stochastic gradient descent to solve the matrix factorization step. We perform 5 repetitions of our experiments for several values of dimension d, T , the parameter \bar{r} of MLingreedy and the regularization parameter $l \in \{l_1, l_2, l_3\}$ of the trace norm policy.

In Figures 3 and 4, we implement MLingreedy with the true value of rank r (most favorable case for this policy in theory) and compare its performance to that of our trace norm bandit and ITL. Trace norm bandit performs consistently above ITL and MLingreedy with true parameter for all considered configurations of T and d . Noticeably MLingreedy with true parameter performs significantly worse than ITL for $d = 40$ for almost all values of N .

In Figure 5, we explore in more details the impact of parameters l and \bar{r} on the performance of the trace norm bandit and MLingreedy, respectively.

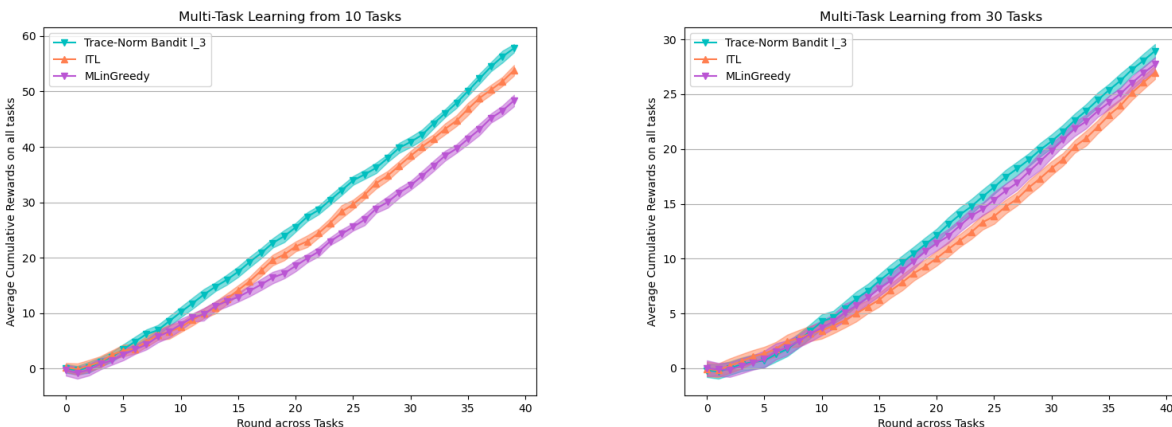


Figure 3: Averaged cumulative reward over all tasks with $d = 10$ features, $T = 10$ tasks (left) and $T = 30$ (right). Each task lasts for $N = 40$ rounds, has $K = 10$ arms, noise variance $\sigma^2 = 1$ and true rank $r = 5$.

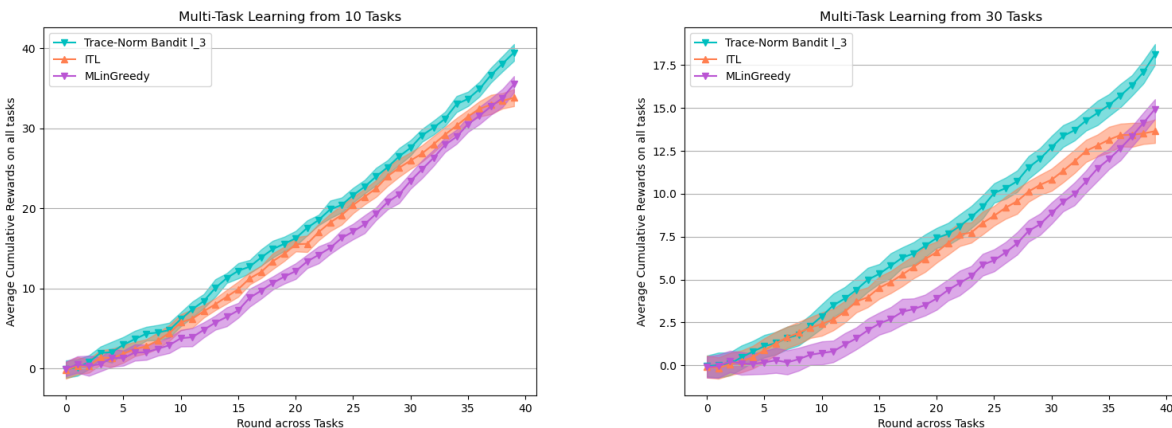


Figure 4: Averaged cumulative reward over all tasks with $d = 40$ features, $T = 10$ tasks (left) and $T = 30$ (right). Each task lasts for $N = 40$ rounds, has $K = 10$ arms, noise variance $\sigma^2 = 1$ and true rank $r = 5$.

- In higher dimension $d = 40$ and $T = 10$, the trace norm bandit performs significantly better than MLinGreedy. Actually MLinGreedy performance is significantly worse than ITL even when using the true rank ($\bar{r} = r$) until N becomes larger than 38 whereas the trace norm policy performs significantly better than ITL even when N is small.
- For $d = 40$ and $T = 30$, the trace norm bandit performs better than ITL for any $N \geq 25$ whereas MLinGreedy with overestimated rank and true rank are below ITL until $N \geq 35$ and $N \geq 38$ respectively. MLinGreedy with underestimated rank is significantly worse than ITL for all N .

In summary, the trace norm bandit performs consistently better than ITL uniformly for most values of N , for all values of d and T and for several choices of the regularization parameter. Conversely, MLinGreedy is far more sensitive to the choice of rank parameter \bar{r} and the values of d and T . MLinGreedy with underestimated rank performs far worse than ITL in all configurations. MLinGreedy with overestimated rank performs sometimes better than MLinGreedy with true rank. Overall, its performance are worse than those of trace norm bandit.

Impact of rank r . In Figure 6, we note that Alg. 1 performs uniformly better than the ITL policy uniformly over all values of rank $r \in [d]$. The performance of the oracle tends to that of ITL as r tends to d . This is expected as the oracle working in a r -dimensional space loses the benefit of working in a small dimensional space when $r \approx d$. Interestingly, the

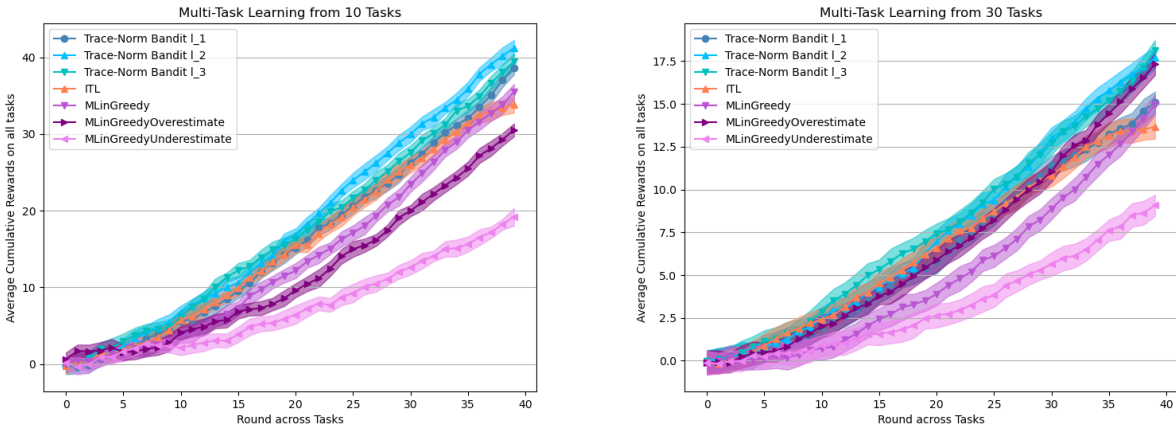


Figure 5: Averaged cumulative reward over all tasks with $d = 40$ features, $T = 10$ tasks (top) and $T = 30$ (bottom). Each task lasts for $N = 40$ rounds, has $K = 10$ arms, noise variance $\sigma^2 = 1$ and true rank $r = 5$.

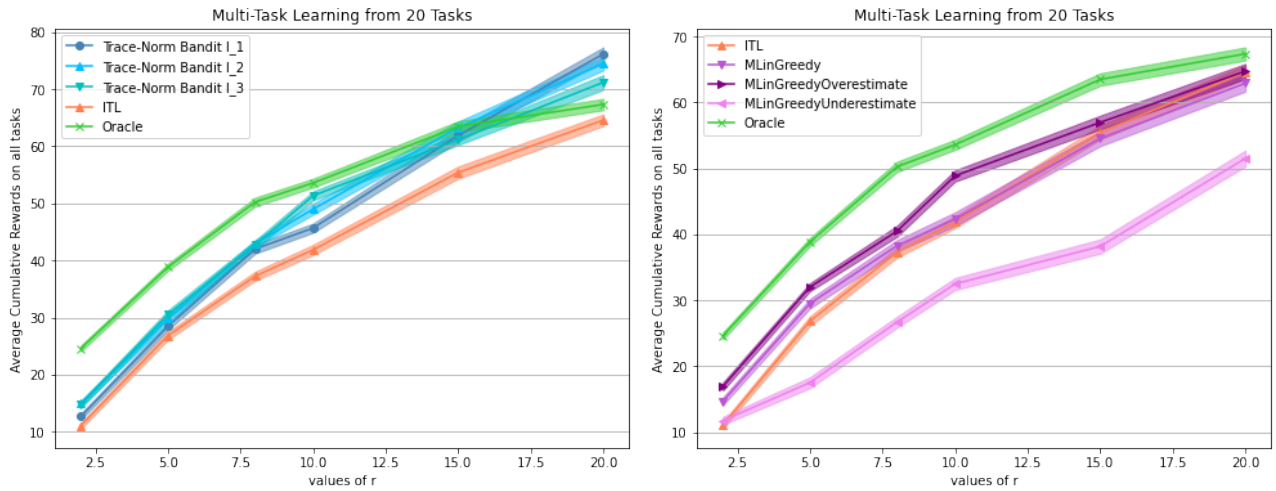


Figure 6: Averaged cumulative reward over all tasks for $d = 20$, $T = 20$, $\sigma^2 = 1$ as a function of the rank r .

performance of the trace norm bandit becomes superior to that of the oracle as r gets close to d . A plausible explanation is that nuclear norm regularization can still bring some benefit even when $r \approx d$ as it can still perform dimension reduction.