
Laplacian-Regularized Graph Bandits: Algorithms and Theoretical Analysis

Kaige Yang
University College London

Xiaowen Dong
University of Oxford

Laura Toni
University College London

Abstract

We consider a stochastic linear bandit problem with multiple users, where the relationship between users is captured by an underlying graph and user preferences are represented as smooth signals on the graph. We introduce a novel bandit algorithm where the smoothness prior is imposed via the random-walk graph Laplacian, which leads to a single-user cumulative regret scaling as $\tilde{\mathcal{O}}(\Psi d\sqrt{T})$ with time horizon T , feature dimensionality d , and the scalar parameter $\Psi \in (0, 1)$ that depends on the graph connectivity. This is an improvement over $\tilde{\mathcal{O}}(d\sqrt{T})$ in **Lin-UCB** [Li et al., 2010], where user relationship is not taken into account. In terms of network regret (sum of cumulative regret over n users), the proposed algorithm leads to a scaling as $\tilde{\mathcal{O}}(\Psi d\sqrt{nT})$, which is a significant improvement over $\tilde{\mathcal{O}}(nd\sqrt{T})$ in the state-of-the-art algorithm **Gob.Lin** [Cesa-Bianchi et al., 2013]. To improve scalability, we further propose a simplified algorithm with a linear computational complexity with respect to the number of users, while maintaining the same regret. Finally, we present a finite-time analysis on the proposed algorithms, and demonstrate their advantage in comparison with state-of-the-art graph-based bandit algorithms on both synthetic and real-world data.

1 Introduction

In the classical multi-armed bandit (MAB) problem, an agent takes sequential actions, choosing one arm

out of the k available ones and it receives an instantaneous payoff from the chosen arm only. The goal of the agent is to learn an action policy that maximizes the cumulative payoff over a course of T rounds [Robbins, 1952]. MAB problems formalize a trade-off between exploration and exploitation, and a particular solution is imposing the principle of optimism in face of uncertainty. Specifically, the agent assigns to each arm an index called the upper confidence bound (UCB) that with high probability is an overestimate of the unknown payoff, and selects the arm with the highest index.

Many variants of the basic MAB problem have been intensively studied, motivated by real-world applications such as ads placement and recommender systems. In the stochastic linear bandit [Auer, 2002], at each round, the agent receives a hint before taking the decision. Specifically, before choosing the arm, the agent is informed of a feature vector $\mathbf{x} \in \mathbb{R}^d$ associated with each arm, referred to as the ‘context’. The payoff associated with each arm is modeled as a noisy linear function of \mathbf{x} with an unknown coefficient vector $\boldsymbol{\theta} \in \mathbb{R}^d$ perturbed by a noise term η , i.e. $y = \mathbf{x}^T \boldsymbol{\theta} + \eta$, where the agent needs to learn $\boldsymbol{\theta}$ based on the context-payoff pair $\{\mathbf{x}, y\}$ and select an arm accordingly. This problem has been well understood in the literature and many studies have already proposed asymptotically optimal algorithms [Auer, 2002, Dani et al., 2008, Agrawal and Goyal, 2013, Chapelle and Li, 2011, Lattimore and Szepesvari, 2016].

The problem is less understood in the case of multiple users, as opposed to a single user, where we assume a central agent needs to select arms for multiple users in a sequential fashion. In this paper, we are interested in the setting where there are n users sharing the same set of arm choices \mathcal{D} containing m arms. The agent faces a set of n independent instance of bandit characterized by an unknown $\boldsymbol{\theta}_i, i \in \{1, 2, \dots, n\}$ specific to each user. At each round, one user out of n is selected uniformly at random, the agent then selects one arm from \mathcal{D} for the user and receives an instantaneous pay-

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

off associated with the selected arm and the user. The overall goal is to minimize the cumulative regret (or equivalently, maximize the cumulative payoffs), which is defined as the sum of instantaneous regret experienced by the agent over a finite time horizon T .

In this setting, naively implementing bandit algorithms on each user independently will result in a cumulative regret that scales linearly with the number of users n . This is clearly infeasible in case of a large number of users. In many cases, however, the users are related in some way and this can be represented by a network (or graph) that encapsulates important additional source of information, such as similarities among users in terms of their preferences (user feature vectors). Exploiting this structure can mitigate the scalability problem. The key setup is therefore to construct a graph where each node represents a user and the edges identify the affinity between users. In real-world applications, such a graph can be a social network of users. This idea leads to a series of work on the so-called graph-based bandit problem [Cesa-Bianchi et al., 2013, Gentile et al., 2014, Liu et al., 2018, Vaswani et al., 2017].

Despite the previous effort, several important limitations still remain to be addressed. First, the graph Laplacian matrix is commonly used in graph-based algorithms, but the justification of its usage remains insufficient (and as to which version of the graph Laplacian leads to optimal policies). As a consequence, the advantage of graph-based bandit is largely shown empirically in previous works, without rigorous theoretical analysis. Furthermore, scalability remains a serious limitation of such algorithms. Involving user graph into bandit algorithms typically results in a computational complexity that scales quadratically with the number of users, which is clearly infeasible in case of large number of users. In this paper, we address the above limitations with the following main contributions:

- We propose a bandit algorithm **GraphUCB** based on the random-walk graph Laplacian, and show its theoretical advantages over other graph Laplacian matrices in reducing cumulative regret. We demonstrate empirically that **GraphUCB** outperforms state-of-the-art graph-based bandit algorithms in terms of cumulative regret.
- As a key ingredient of the proposed algorithm, we derive a novel UCB representing the single-user bound while embedding the graph structure, which reduces the size of the confidence set, in turn leading to lower regret;
- To improve scalability, we further propose a simplified algorithm **GraphUCB-Local** whose complexity scales linearly with respect to the number

of users, yet still holding the same regret upper bound as **GraphUCB**;

- Finally, we derive a finite-time analysis on both algorithms and show a lower regret upper bound than other state-of-the-art graph-based bandit algorithms.

2 Related work

Graph-based bandit algorithms can be roughly categorized as: *i*) topology-based bandits, where the graph topology itself is exploited to improve learning performance, and *ii*) spectral bandits, where the user feature vectors θ are modeled as signals defined on the underlying graph, whose characteristics are then exploited in the graph spectral domain via tools provided by graph signal processing [Shuman et al., 2013] to assist learning.

In topology-based bandits, the key intuition is to achieve dimensionality reduction in the user space by exploiting the graph topology. Specifically, users can be clustered based on the graph topology and a per-cluster feature vector can be learned, substantially reducing the dimensionality of the problem as opposed to the case in which one vector is learned per user. For example, [Gentile et al., 2014] clusters users based on the connected components of the user graph, and [Li et al., 2016] generalizes it to consider both the user graph and item graph. On the other hand, [Yang and Toni, 2018] makes use of community detection techniques on graphs to find user clusters. More broadly, in the spirit of dimensionality reduction, even without constructing an explicit user graph, the work in [Korda et al., 2016] proposes a distributed clustering algorithm while [Nguyen and Lauw, 2014] applies *k-means* clustering to the user features. Despite the differences in the proposed techniques, these studies share two common drawbacks: 1) the learning performance depends on the clustering algorithm being used, which tends to be expensive for large-scale graphs; 2) learning a per-cluster (and not per-user) feature vector means ignoring the subtle difference between users within the same cluster. In short, clustering can reduce the dimensionality of the user space, but it does not necessarily preserve key users characteristics. To achieve both goals simultaneously, there is a need for a proper mathematical framework able to incorporate the user relationship into learning in a more direct way.

On the spectral bandit side, the strong assumption that users can be grouped into clusters is relaxed; users are assumed to be similar with their neighbors in the graph and such similarity is reflected by the weight of graph edges. [Wu et al., 2016] employs a graph

Laplacian-regularized estimator, which promotes similar feature vectors for users connected in the graph. In their setting, however, each arm is selected by all users jointly. This work results in a network regret scaling with $\tilde{O}(dn\sqrt{T})$. [Vaswani et al., 2017] casts the same estimator as GMRF (Gaussian Markov Random Field) and proposes a Thompson sampling algorithm, leading to a much simpler algorithmic implementation without a UCB evaluation. However, the regret bound remains $\tilde{O}(dn\sqrt{T})$. To the best of our knowledge, an efficient algorithm able to address the multi-user MAB problem with a sub-linear regret bound is still missing.

A proper bound and mathematical derivation of spectral MAB are provided in [Valko et al., 2014], which represents the payoffs of arms as smooth signals on a graph with the arms being the nodes. Specifically, the arm features \mathbf{x} are modeled as eigenvectors of the graph Laplacian and the sparsity of such eigenvectors is exploited to reduce the dimensionality of \mathbf{x} , to a so-called ‘effective dimension’ term \tilde{d} . This work shows an improved regret bound $\tilde{O}(\tilde{d}\sqrt{T})$ where \tilde{d} is significant less than d in **LinUCB** [Abbasi-Yadkori et al., 2011]. While interesting, the proposed solution applies to the single-user with high-dimensional arm set. Whereas, in our setting, the dimensionality issue is caused by the large number of users, leading to a completely different mathematical problem.

Among these works on graph bandit, the one that is most similar to our work in terms of problem definition and proposed solution is [Cesa-Bianchi et al., 2013]. In [Cesa-Bianchi et al., 2013], the graph is exploited such that each user shares instantaneous payoff with neighbors, which is promoted by a Laplacian-regularized estimator. This implicitly imposes smoothness among the feature vectors of users, resulting in the estimate of similar feature vectors for users connected by edges with strong weights in the graph.

In our paper, we proposed the **GraphUCB** algorithm that builds on and improves the work of **Gob.Lin** in a number of important ways:

- First, **Gob.Lin** employs the combinatorial Laplacian as a regularizer, whereas our algorithm **GraphUCB** makes use of the random-walk graph Laplacian. We prove theoretically that the combinatorial Laplacian results in a cumulative regret scaling with the number of users, which could be large. However, random-walk graph Laplacian overcomes this serious drawback and yields a sub-linear regret with the number of users.
- Second, UCB used in **Gob.Lin** results in a cumulative regret that scales with $\tilde{O}(nd\sqrt{T})$ which is worse than **LinUCB** $\tilde{O}(d\sqrt{nT})$. We propose a new UCB that leads to a cumulative regret scaling

with $\tilde{O}(\Psi d\sqrt{nT})$ where $\Psi \in (0, 1)$.

- Finally, the computational complexity of **Gob.Lin** is quadratic with respect to the number of users. Our simplified algorithm **GraphUCB-Local** scales linearly with the number of users, and at the same time enjoys the same regret bound as **GraphUCB**. This significantly improves the scalability of the proposed graph bandit algorithm.

3 Setting

We consider a linear bandit problem with m arms and n users. We denote by \mathcal{U} the user set with cardinality $|\mathcal{U}| = n$ and by \mathcal{D} the arm set with $|\mathcal{D}| = m$. Each arm is described by a feature vector $\mathbf{x} \in \mathbb{R}^d$, while each user is described by a parameter vector $\boldsymbol{\theta} \in \mathbb{R}^d$, with d being the dimension of both vectors. The affinity between users is encoded by an undirected and weighted graph $\mathcal{G} = (V, E)$, where $V = \{1, 2, \dots, n\}$ represents the node set for n users and E represents the edge set. The graph \mathcal{G} is known a priori and identified by its adjacency matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$, where $W_{ij} = W_{ji}$ captures the affinity between $\boldsymbol{\theta}_i$ and $\boldsymbol{\theta}_j$. The combinatorial Laplacian of \mathcal{G} is defined as $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where \mathbf{D} is a diagonal matrix with $D_{ii} = \sum_{i=1}^n W_{ii}$. The symmetric normalized Laplacian is defined as $\tilde{\mathbf{L}} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$. In addition, the random-walk graph Laplacian is defined as $L = \mathbf{D}^{-1} \mathbf{L}$.

In our setting, the unknown user features $\boldsymbol{\Theta} = [\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_n]^T \in \mathbb{R}^{n \times d}$ are assumed to be smooth over \mathcal{G} . The smoothness of $\boldsymbol{\Theta}$ over graph \mathcal{G} can then be quantified using the Laplacian quadratic form of any of the three Laplacian defined above. In this work, we choose the random-walk graph Laplacian L because of its two unique properties $\mathcal{L}_{ii} = 1$ and $\sum_{j \in i} \mathcal{L}_{ij} = -1$. The benefit of these properties will be clear after the introduction of our proposed bandit algorithm (see Remark 1). Mathematically, the Laplacian quadratic form based on L is (see Appendix H for the derivation):

$$tr(\boldsymbol{\Theta}^T L \boldsymbol{\Theta}) = \frac{1}{4} \sum_{k=1}^d \sum_{i,j} \left(\frac{W_{ij}}{D_{ii}} + \frac{W_{ji}}{D_{jj}} \right) (\Theta_{ik} - \Theta_{jk})^2 \quad (1)$$

where Θ_{ik} is the (i, k) -th element of $\boldsymbol{\Theta}$. The more the graph \mathcal{G} reflects the similarity between users correctly, the smaller the quadratic term $tr(\boldsymbol{\Theta}^T L \boldsymbol{\Theta})$. Specifically, $tr(\boldsymbol{\Theta}^T L \boldsymbol{\Theta})$ is small when Θ_{ik} and Θ_{jk} are similar given a large weight $\frac{W_{ij}}{D_{ii}} + \frac{W_{ji}}{D_{jj}}$.

Equipped with the above notation, we now introduce the multi-user bandit problem, in which an agent needs to take sequential decisions (e.g., recommendations) for a multitude of users appearing over time. At each

time $t = 1, \dots, T$, the agent is informed about the user i_t to serve, with the user being selected uniformly at random from the user set \mathcal{U} . Then, the agent selects an arm $\mathbf{x}_t \in \mathcal{D}$ to be recommended to user i_t . Upon this selection, the agent observes a payoff y_t , which is assumed to be generated by noisy versions of linear functions of the users and item vectors. Namely,

$$y_t = \mathbf{x}_t^T \boldsymbol{\theta}_{i_t} + \eta_t \quad (2)$$

where the noise η_t is assumed to be σ -sub-Gaussian for any t .

The agent is informed about the graph \mathcal{G} and the arm feature vectors \mathbf{x}_a , $a \in \{1, 2, \dots, m\}$, while $\boldsymbol{\Theta}$ is unknown and needs to be inferred. The goal of the agent is to learn a selection strategy that minimizes the cumulative regret with respect to an optimal strategy, which always selects the optimal arm for each user. Formally, after a time horizon T , the cumulative (pseudo) regret is defined as:

$$R_T = \sum_{t=1}^T (\mathbf{x}_t)^T \boldsymbol{\theta}_{i_t} - \mathbf{x}_t^* T \boldsymbol{\theta}_{i_t} \quad (3)$$

where \mathbf{x}_t and \mathbf{x}_t^* are the arm selected by the agent and the optimal strategy at t , respectively. Note that the optimal choice depends on t as well as on the user i_t . For notation convenience in the rest of the paper, at each time step t , we use i to generally refer to the user appeared and \mathbf{x}_t to represent the feature vector of the arm selected.

4 Laplacian-Regularized Estimator

To estimate the user parameter $\boldsymbol{\Theta}$ at time t , we make use of the Laplacian-regularized estimator:

$$\hat{\boldsymbol{\Theta}}_t = \arg \min_{\boldsymbol{\Theta} \in \mathbb{R}^{n \times d}} \sum_{i=1}^n \sum_{\tau \in \mathcal{T}_{i,t}} (\mathbf{x}_\tau^T \boldsymbol{\theta}_i - y_\tau)^2 + \alpha \text{tr}(\boldsymbol{\Theta}^T \mathbf{L} \boldsymbol{\Theta}) \quad (4)$$

where $\boldsymbol{\theta}_i$ is the i -th row of $\boldsymbol{\Theta}$, $\mathcal{T}_{i,t}$ is the set of time steps at which user i is served up to time t . \mathbf{x}_τ is the feature of arm selected by the learner, y_τ is the payoff from user i at time τ , and α is the regularization parameter. Clearly, Eq. 4 is convex and can be solved via convex optimization techniques. Specifically, it has a closed form solution [Alvarez et al., 2012]:

$$\text{vec}(\hat{\boldsymbol{\Theta}}_t) = (\boldsymbol{\Phi}_t \boldsymbol{\Phi}_t^T + \alpha \mathbf{L} \otimes \mathbf{I})^{-1} \boldsymbol{\Phi}_t \mathbf{Y}_t \quad (5)$$

where \otimes is the Kronecker product, $\text{vec}(\hat{\boldsymbol{\Theta}}_t) \in \mathbb{R}^{nd}$ is the concatenation of the columns of $\hat{\boldsymbol{\Theta}}_t$, $\mathbf{I} \in \mathbb{R}^d \times \mathbb{R}^d$ is the identity matrix, and $\mathbf{Y}_t = [y_1, y_2, \dots, y_t]^T \in \mathbb{R}^t$ is the collection of all payoffs. Finally, $\boldsymbol{\Phi}_t = [\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \dots, \boldsymbol{\phi}_t] \in \mathbb{R}^{nd \times t}$, where $\boldsymbol{\phi}_t \in \mathbb{R}^{nd}$, is a long

sparse vector indicating that the arm with feature \mathbf{x}_t is selected for user i . Formally,

$$\boldsymbol{\phi}_t^T = \left(\begin{array}{c|c} \mathbf{0}_{(i-1) \times d} & \mathbf{x}_t^T \\ \hline \mathbf{0}_{(n-i) \times d} & \mathbf{0}_{1 \times d} \end{array} \right). \quad (6)$$

While Eq. 4 provides the estimate of feature vectors of all users at t , i.e., $\hat{\boldsymbol{\Theta}}_t$, the agent is interested in the estimation of each single-user feature vector $\hat{\boldsymbol{\theta}}_{i,t}$. Mathematically, $\hat{\boldsymbol{\theta}}_{i,t}$ can be obtained by decoupling users in Eq. 5. This however is highly complex due to the inversion $(\boldsymbol{\Phi}_t \boldsymbol{\Phi}_t^T + \alpha \mathbf{L} \otimes \mathbf{I})^{-1}$, which the agent needs to preform at each time step (we recall that the Laplacian is high-dimensional). We notice that the close-form solution of $\hat{\boldsymbol{\theta}}_{i,t}$ can be closely approximated via a Taylor expansion of $(\boldsymbol{\Phi}_t \boldsymbol{\Phi}_t^T + \alpha \mathbf{L} \otimes \mathbf{I})^{-1}$, as stated in Lemma 1 and further commented and tested empirically in Appendix A.

Lemma 1. $\hat{\boldsymbol{\theta}}_{i,t}$ is obtained from Eq. 5, let $\hat{\boldsymbol{\theta}}_{i,t}$ be the i -th row of $\hat{\boldsymbol{\Theta}}_t$ which is the estimate of $\boldsymbol{\theta}_i$. $\hat{\boldsymbol{\theta}}_{i,t}$ can be approximated by¹:

$$\hat{\boldsymbol{\theta}}_{i,t} \approx \mathbf{A}_{i,t}^{-1} \mathbf{X}_{i,t} \mathbf{Y}_{i,t} - \alpha \mathbf{A}_{i,t}^{-1} \sum_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \mathbf{X}_{j,t} \mathbf{Y}_{j,t} \quad (7)$$

where $\mathbf{A}_{i,t} = \sum_{\tau \in \mathcal{T}_{i,t}} \mathbf{x}_\tau \mathbf{x}_\tau^T \in \mathbb{R}^{d \times d}$ is the Gram matrix related to the choices made by user i , $\mathcal{T}_{i,t}$ is the set of time at which user i is served up to time t , and \mathcal{L}_{ij} is the (i, j) -th element in \mathcal{L} . $\mathbf{X}_{i,t} \in \mathbb{R}^{d \times |\mathcal{T}_{i,t}|}$ is the collection of features of arms that are selected for user i up to time t with $\{\mathbf{x}_\tau\}, \tau \in \mathcal{T}_{i,t}$ as columns. $\mathbf{Y}_{i,t} \in \mathbb{R}^{|\mathcal{T}_{i,t}|}$ is the collection of payoffs associated with user i up to time t , whose elements are $\{y_\tau\}, \tau \in \mathcal{T}_{i,t}$.

Proof. See Appendix A. \square

4.1 Construction of Confidence Set

For the agent to balance exploration and exploitation in sequential decisions, we need to quantify the uncertainty over the estimation of $\hat{\boldsymbol{\theta}}_{i,t}$. This is possible by defining a confidence set around $\hat{\boldsymbol{\theta}}_{i,t}$ based on Mahalanobis distance using its precision matrix, as commonly adopted in bandit literature [Lattimore and Szepesvári, 2018]. Let $\mathbf{A}_{i,t} \in \mathbb{R}^{d \times d}$ be the precision matrix of $\hat{\boldsymbol{\theta}}_{i,t}$, the confidence set is formally defined as:

$$\mathcal{C}_{i,t} = \{\boldsymbol{\theta}_{i,t} : \|\hat{\boldsymbol{\theta}}_{i,t} - \boldsymbol{\theta}_{i,t}\|_{\mathbf{A}_{i,t}} \leq \beta_{i,t}\} \quad (8)$$

where $\beta_{i,t}$ is the upper bound of $\|\hat{\boldsymbol{\theta}}_{i,t} - \boldsymbol{\theta}_{i,t}\|_{\mathbf{A}_{i,t}}$ which is what we are interested in for the bandit algorithm.

¹ $\mathbf{A}_{i,t}$ (and $\mathbf{A}_{j,t}$) is not full-rank when $|\mathcal{T}_{i,t}| < d$. To guarantee inversion, in practice we set $\mathbf{A}_{i,t} = \sum_{\tau \in \mathcal{T}_{i,t}} \mathbf{x}_\tau \mathbf{x}_\tau^T + \mathbf{I}_d$ with $\epsilon = 0.01$.

With this goal in mind, we seek an expression for $\mathbf{\Lambda}_{i,t}$. Let $\mathbf{\Lambda}_t \in \mathbb{R}^{nd \times nd}$ denote the precision matrix of $\text{vec}(\hat{\Theta}_t) \in \mathbb{R}^{nd}$, where $\mathbf{\Lambda}_{i,t} \in \mathbb{R}^{d \times d}$ is the i -th block matrix along the diagonal of $\mathbf{\Lambda}_t$. Defining the precision matrix of $\text{vec}(\hat{\Theta}_t) \in \mathbb{R}^{nd}$ as:

$$\mathbf{\Lambda}_t = \mathbf{M}_t \mathbf{A}_t^{-1} \mathbf{M}_t \quad (9)$$

with $\mathbf{A}_t = \Phi_t \Phi_t^T$, $L = L \mathbf{I}$, and $\mathbf{M}_t = \mathbf{A}_t + \alpha L$, we have:

$$\mathbf{\Lambda}_{i,t} = \mathbf{A}_{i,t} + 2\alpha \mathcal{L}_{ii} \mathbf{I} + \alpha^2 \sum_{j=1}^n \mathcal{L}_{ij}^2 \mathbf{A}_{j,t}^{-1} \quad (10)$$

where $\mathbf{A}_{i,t}$ and $\mathbf{A}_{j,t}$ are defined in Lemma 1, and \mathcal{L}_{ij} is the (i, j) -th element in L . A detailed derivation of Eq. 10 is presented in Appendix B and Appendix C. Given Eq. 10, we can upper bound the size of the confidence set, which provides the value of $\beta_{i,t}$.

Lemma 2. *Let $\mathbf{V}_{i,t} = \mathbf{A}_{i,t} + \alpha \mathcal{L}_{ii} \mathbf{I}$, and $\mathbf{I} \in \mathbb{R}^{d \times d}$ is the identity matrix. Given a scalar $\delta \in [0, 1]$, and by defining $\mathbf{\Delta}_i = \sum_{j=1}^n \mathcal{L}_{ij} \theta_j$, the size of the confidence set defined in Eq. 8 is upper bounded with probability $1 - \delta$ by $\beta_{i,t}$:*

$$\beta_{i,t} = \sigma \sqrt{2 \log \frac{|\mathbf{V}_{i,t}|^{1/2}}{\delta |\alpha \mathbf{I}|^{1/2}} + \sqrt{\alpha} \|\mathbf{\Delta}_i\|_2} \quad (11)$$

Proof. See Appendix D. \square

It is worth mentioning that the bound $\beta_{i,t}$ depends on $\mathbf{\Delta}_i$, which reflects information contained in the graph structure. In the case of random-walk Laplacian in which $\mathcal{L}_{ii} = 1$ and $\sum_{j \neq i} \mathcal{L}_{ij} = 1$, $\mathbf{\Delta}_i = \theta_i - \sum_{j \neq i} \mathcal{L}_{ij} \theta_j$ measures the difference between user feature θ_i and the weighted average of its neighbors. To show the effect of graph structure on θ_i , we consider two extreme cases: *a)* an empty graph², i.e., $\mathcal{L}_{ii} = 1$ and $\mathcal{L}_{ij} = 0$. In this case, $\mathbf{\Delta}_i = \theta_i$, which recovers the confidence set used in **LinUCB** [Abbasi-Yadkori et al., 2011]; *b)* a fully connected graph with $W_{ij} = 1$ which means $\mathcal{L}_{ii} = 1$, $\mathcal{L}_{ij} = \frac{1}{n-1}$ and $\theta_i = \theta_j$. In this case, $\mathbf{\Delta}_i = \theta_i - \frac{n-1}{n-1} \theta_i = \mathbf{0}$, leading to a much lower bound than the one in case *a)* and with no graph structure to be exploited. In between, $\mathbf{\Delta}_i$ depends on the similarity between θ_i and its neighbors $\theta_j, j \neq i$. In general, the smoother the signal is on the graph (in the sense of a small Laplacian quadratic in Eq. 1), the lower the $\|\mathbf{\Delta}_i\|_2$. This has been empirically shown in Figure 1(a), where we depict $\|\mathbf{\Delta}_i\|_2$ as a function of the level of smoothness quantified by $\text{tr}(\Theta^T L \Theta)$ where smaller value means smoother Θ over the graph.

²For isolated node, we set $L_{ii} = 1$.

Algorithm 1: GraphUCB

Input : α, T, L, δ

Initialization : For any $i \in \{1, 2, \dots, n\}$

$$\hat{\theta}_{0,i} = \mathbf{0} \in \mathbb{R}^d, \mathbf{\Lambda}_{0,i} = \mathbf{0} \in \mathbb{R}^{d \times d},$$

$$\mathbf{A}_{0,i} = \mathbf{0} \in \mathbb{R}^{d \times d}, \beta_{i,t} = 0.$$

for $t \in [1, T]$ **do**

 User index i is selected

$$1. \mathbf{A}_{i,t} \leftarrow \mathbf{A}_{i,t-1} + \mathbf{x}_t \mathbf{x}_t^T$$

$$2. \mathbf{A}_{j,t} \leftarrow \mathbf{A}_{j,t-1}, \forall j \neq i.$$

3. Update $\mathbf{\Lambda}_{i,t}$ via Eq. 10.

4. Select \mathbf{x}_t via Eq. 12

 where $\beta_{i,t}$ is defined in Eq. 11

5. Receive the payoff y_t

6. Update $\hat{\Theta}_t$ via Eq. 4

end

Now that we have introduced $\mathbf{\Delta}_i$, we are ready to motivate the choice of the random-walk graph Laplacian instead of other commonly used graph Laplacians.

Remark 1. *The two unique properties $\mathcal{L}_{ii} = 1$ and $\sum_{j \neq i} \mathcal{L}_{ij} = 1$ of the random-walk graph Laplacian L ensure a bounded regret and lower regret with more similar users. The same cannot be guaranteed with the combinatorial or normalized Laplacian.*

To see this more clearly, if the combinatorial Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W}$ is used in the Laplacian-regularized estimator (Eq. 4), the term $\mathbf{\Delta}_i$ becomes $\mathbf{\Delta}_i = D_{ii}(\theta_i + \sum_{j \neq i} \frac{W_{ij}}{D_{ii}} \theta_j)$. This term scales linearly with D_{ii} , the degree of each user i , resulting in a regret that also scales with D_{ii} and may become rather large for densely connected graphs. On the other hand, if the symmetric normalized Laplacian $\tilde{\mathbf{L}} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$ is used in Eq. 4, we have $\sum_{j \neq i} \tilde{L}_{ij} \neq 1$. It follows that $\sum_{j \neq i} \tilde{L}_{ij} \theta_j$ will not be a convex combination of $\theta_j, j \neq i$, which means that there is no guarantee for $\mathbf{\Delta}_i$ to be located inside the Euclidean ball defined by $\|\theta_i\|_2 \leq 1, \forall i \in [1, \dots, n]$, leading to an unbounded regret. By contrast, the two unique properties $\mathcal{L}_{ii} = 1$ and $\sum_{j \neq i} \mathcal{L}_{ij} = 1$ of the random-walk normalized Laplacian L ensure a bounded regret and less regret if users are similar.

5 Algorithms

We now introduce our proposed **GraphUCB** bandit algorithm, sketched in **Algorithm 1** and based on the principle of *optimism in face of uncertainty*:

$$\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{D}} \mathbf{x}^T \hat{\theta}_{i,t} + \beta_{i,t} \|\mathbf{x}\|_{\mathbf{\Lambda}_{i,t}^{-1}} \quad (12)$$

GraphUCB is designed based on the Laplacian regularized estimator Eq. 4 and the arm selection principle

in Eq. 12. Formally, at each time t , an user index i is selected randomly from the user set \mathcal{U} . The algorithm first updates $\mathbf{\Lambda}_{i,t}$ based on the Gram matrix $\mathbf{A}_{i,t}$ and $\mathbf{A}_{j,t}$. Then, it selects the arm \mathbf{x}_t from the arm set \mathcal{D} following Eq. 12. Upon receiving the instantaneous payoff y_t , it updates the features of all users $\hat{\Theta}_t$ by solving Eq. 4. The process then continues to time $t+1$, and is repeated until T . In Eq. 11 the Δ_i is based on the unknown ground-truth θ_i and θ_j . In practice, it is replaced by its empirical counterpart $\hat{\Delta}_i = \frac{1}{n} \sum_{j=1}^n \mathcal{L}_{ij} \hat{\theta}_{i,t}$ where $\hat{\theta}_{i,t}$ is the i -th row of $\hat{\Theta}_t$.

One limitation of **GraphUCB** is its high computational complexity. Specifically, in solving Eq. 4, the running time is dominated by the inversion $(\Phi_t \Phi_t^T + \alpha L \otimes \mathbf{I})^{-1}$, which is in the order $\mathcal{O}(n^2 d^2)$. This is impractical when the user number n is large. Recall that in the learning setting, at each time t , only one user is selected. Thus, it suffices to only update $\hat{\theta}_{i,t}$ (i.e., a local rather than global update). Therefore, we propose to make use of Lemma 1 instead of Eq. 4 to estimate $\hat{\theta}_{i,t}$, which results in a significant reduction in computational complexity. Clearly, the complexity of the approximation in Lemma 1 is dominated by the inversion of $(\mathbf{A}_{i,t} + \alpha \mathcal{L}_{ii} \mathbf{I})^{-1}$ or $\mathbf{A}_{j,t}^{-1}$, which is in the order $\mathcal{O}(d^2)$. Since the approximation involves n such inversions, the total complexity is $\mathcal{O}(n d^2)$, i.e., it scales linearly (rather than quadratically) with n .

By using Lemma 1, we therefore propose a second algorithm **GraphUCB-Local** serving as a simplified version of **GraphUCB**. The only difference lies in the number of users updated per round. **GraphUCB** updates all users $\hat{\Theta}_t$ via Eq. 4 (closed-form solution), while **GraphUCB-Local** only updates one user $\hat{\theta}_{i,t}$ via Lemma 1 (approximated solution). Pseudocode of **GraphUCB-Local** is presented in Appendix E.

6 Analysis

Before providing the finite-time analysis on the proposed algorithms, we define:

$$\Psi_{i,T} = \frac{\mathbb{P} \left[\prod_{\tau \in \mathcal{T}_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2 \right]}{\mathbb{P} \left[\prod_{\tau \in \mathcal{T}_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2 \right]} \quad (13)$$

where $\mathcal{T}_{i,T}$ is the set of time steps in which user i is served up to time T , $\mathbf{A}_{i,\tau} = \sum_{\ell \in \mathcal{T}_{i,\tau}} \mathbf{x}_\ell \mathbf{x}_\ell^T$, $\mathbf{V}_{i,\tau} = \mathbf{A}_{i,\tau} + \alpha \mathcal{L}_{ii} \mathbf{I}$ and $\mathbf{\Lambda}_{i,\tau}$ is defined as in Eq. 10. Moreover, $\|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2$ and $\|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2$ quantify the variance of

predicted payoff $\hat{y}_\tau = \hat{\theta}_{i,\tau}^T \mathbf{x}_\tau$ in the cases where the graph structured is exploited or ignored, respectively.

Lemma 3. Let $\Psi_{i,T}$ be as defined in Eq. 13 and $\|\mathbf{x}_\tau\|_2 \leq 1$ for any $\tau \leq T$, then:

$$\Psi_{i,T} \in (0, 1)$$

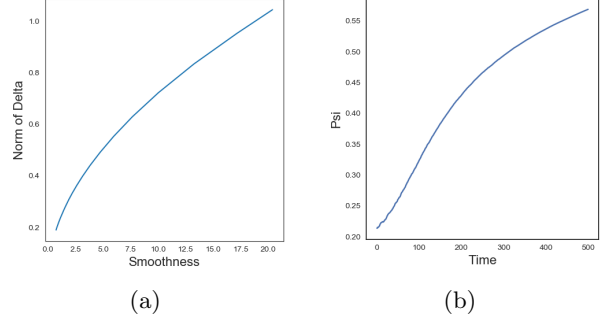


Figure 1: (a) $\|\Delta_i\|_2$ vs. smoothness ($\text{tr}(\Theta^T L \Theta)$), (b) $\Psi_{i,T}$ vs. time.

and as $T \rightarrow \infty$, $\Psi_{i,T} \rightarrow 1$. This implies:

$$\prod_{\tau \in \mathcal{T}_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2 \leq \prod_{\tau \in \mathcal{T}_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2.$$

Proof on the bound $\Psi_{i,T}$ is provided in Appendix F. The Lemma highlights the importance in taking into account the graph structure in the payoff estimation, showing that the uncertainty of \hat{y}_τ is reduced when the graph structure is exploited. This effect diminishes with time: in Fig. 1(b), we see that as more data are collected, the graph-based estimator approaches the estimator in which users parameters are estimated independently (since $\Psi_{i,T} \rightarrow 1$).

6.1 Regret Upper Bound

We present the cumulative regret upper bounds satisfied by both **GraphUCB** and **GraphUCB-Local**.

Theorem 1. $\Psi_{i,T}$ is defined in Eq. 13, $\mathbf{\Lambda}_{i,T}$ defined in Eq. 10 and $\Delta_i = \frac{1}{n} \sum_{j=1}^n \mathcal{L}_{ij} \theta_j$. Without loss of generality, assume $\|\theta_i\|_2 \leq 1$ for any $i \in \{1, 2, \dots, n\}$ and $\|\mathbf{x}_\tau\|_2 \leq 1$ for any $\tau \leq T$. Then, for $\delta \in [0, 1]$, for any user $i \in \{1, 2, \dots, n\}$ the cumulative regret over time horizon T satisfies the following upper bound with probability $1 - \delta$:

$$R_{i,T} = \sum_{\tau \in \mathcal{T}_{i,T}} r_\tau = \mathcal{O} \left(\frac{1}{\Psi_{i,T}} \frac{1}{d \log(|\mathcal{T}_{i,T}|) + \sqrt{\alpha} \|\Delta_i\|_2} \right) \times \Psi_{i,T} \frac{1}{d |\mathcal{T}_{i,T}| \log(|\mathcal{T}_{i,T}|)} \quad (14)$$

Assuming that users are served uniformly up to time horizon T , i.e., $|\mathcal{T}_{i,T}| = T/n$, the network regret (the total cumulative regret experienced by all users) satisfies the following upper bound with probability $1 - \delta$:

$$R_T = \sum_{i=1}^n R_{i,T} = \sum_{i=1}^n \tilde{\mathcal{O}} \left(\frac{1}{\Psi_{i,T}} \frac{1}{d \sqrt{T/n}} \right) = \tilde{\mathcal{O}} \left(d \sqrt{Tn} \max_{i \in \mathcal{U}} \Psi_{i,T} \right) \quad (15)$$

Proof. See Appendix G. \square

Remark 2. Both *GraphUCB* and *GraphUCB-Local* satisfy Theorem 1.

The regret upper bound in Theorem 1 is derived based on **GraphUCB-Local** algorithm (Appendix G). Due to the approximation error introduced in the Taylor expansion, the regret of **GraphUCB-Local** is worse than that of **GraphUCB**. Therefore, Theorem 1 is also an upper bound of **GraphUCB**.

6.2 Comparison with LinUCB and Gob.Lin

Under the same setting, the single-user regret upper bound of **LinUCB** [Li et al., 2010] is:

$$R_{i,T} = \mathcal{O} \left(\frac{\sqrt{d \log(|\mathcal{T}_{i,T}|)} + \sqrt{\alpha} \|\boldsymbol{\theta}_i\|_2}{d |\mathcal{T}_{i,T}| \log(|\mathcal{T}_{i,T}|)} \right) \quad (16)$$

Since $\|\boldsymbol{\Delta}_i\|_2 \leq \|\boldsymbol{\theta}_i\|_2$ and $\Psi_{i,T} \in (0,1)$ (Lemma 2 and Lemma 3), **GraphUCB** (and **GraphUCB-Local**) leads to a lower regret –Eq. 14– than **LinUCB** –Eq. 16.

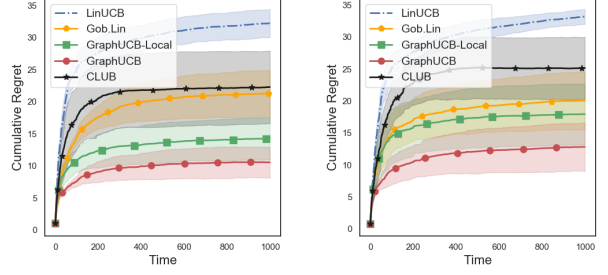
The cumulative regret experienced by all users in **Gob.Lin** [Cesa-Bianchi et al., 2013] is upper bounded by:

$$R_T = 4 \sqrt{T(\sigma^2 \log \frac{|\mathbf{M}_t|}{\delta} + L(\boldsymbol{\theta}) \log |\mathbf{M}_t|)} = \tilde{\mathcal{O}}(nd\sqrt{T}) \quad (17)$$

where $L(\boldsymbol{\theta}) = \sum_{i=1}^n \|\boldsymbol{\theta}_i\|_2 + \sum_{(i,j) \in E} \|\boldsymbol{\theta}_i - \boldsymbol{\theta}_j\|_2$. Clearly, the cumulative regret achieved by **GraphUCB** in Eq. 15 is in the order of $\tilde{\mathcal{O}}(\sqrt{n})$ which is better than that in Eq. 17. This is mainly due to the different UCBs used in these algorithms. Specifically, **Gob.Lin** proposed the following bound:

$$\beta_t = \sigma \sqrt{\log \frac{|\mathbf{M}_t|}{\delta} + L(\boldsymbol{\theta})} = \tilde{\mathcal{O}}(\sqrt{nd}), \quad (18)$$

while we propose a lower single-user bound $\beta_{i,t} = \tilde{\mathcal{O}}(\sqrt{d})$ (Eq. 11). As described in Remark 1, the bound in Eq. 18 grows with the degree of the network, which in the bound is hidden in $L(\boldsymbol{\theta})$. We emphasize that in practice **GOB.Lin** could perform much better than its regret upper bound Eq. 17, if β_t defined in Eq. 18 is replaced by $\lambda \sqrt{\log(t+1)}$ where λ is a tunable parameter. This is exactly what the authors of **Gob.Lin** did in their original paper when reporting empirical results. We follow the same trick when implementing **Gob.Lin** in our experiments, and show that **GraphUCB** still achieves better performance empirically.



(a) RBF ($s=0.5$)

(b) ER ($p=0.4$)

Figure 2: Cumulative regret vs. time for different type of graphs (ER and RBF) consistently generated with the same level of smoothness and sparsity between graphs.

7 Experiment Results

We evaluate the proposed algorithms and compare them to **LinUCB** (no graph information exploited in the bandit), **Gob.Lin** (graph exploited in the features estimation) and **CLUB** (graph exploited to cluster users). All results reported are averaged across 20 runs. In all experiments, we set confidence probability parameter $\delta = 0.01$, noise variance $\sigma = 0.01$, and regularization parameter $\alpha = 1$. For **Gob.Lin**, we use $\beta_{i,t} = \lambda \sqrt{\log(t+1)}$, and λ is set using the best value in range $[0, 1]$. For **CLUB**, the edge deletion parameter α_2 is tuned to its best value.

7.1 Experiments on Synthetic Data

In the synthetic simulations, we first generate a graph \mathcal{G} and then generate a smooth Θ via Eq. 19 which is proposed in [Yankelevsky and Elad, 2016]:

$$\Theta = \arg \min_{\mathbb{R}^{n \times d}} \|\Theta - \Theta_0\|_F^2 + \gamma \text{tr}(\Theta^T \mathcal{L} \Theta) \quad (19)$$

where $\Theta_0 \in \mathbb{R}^{n \times d}$ is a randomly initialized matrix, and L is the random-walk graph Laplacian of \mathcal{G} . The second term in Eq. 19 promotes the smoothness of Θ : the larger the γ , the smoother the Θ over the graph³. In all experiments, $n = 20, d = 5$. To simulate \mathcal{G} , we follow two random graph models commonly used in the network science community: 1) Radial basis function (RBF) model, a weighted fully connected graph, with edge weights $W_{ij} = \exp(-\rho \|\boldsymbol{\theta}_i - \boldsymbol{\theta}_j\|^2)$; 2) Erdős Rényi (ER) model, an unweighted graph, in which each edge is generated independently and randomly with probability p .

³The regularization parameter in Eq. 19 is used to generate a smooth function in the synthetic settings, while the parameter in Eq. 4 is used in the bandit algorithm to infer the smooth prior when estimating user features.




Figure 3: Cumulative regret for RBF graphs with different level of smoothness (a) and sparsity (b).

In Figure 2, we depict the cumulative per-user regret as a function of time for both RBF and ER graphs for both our proposed algorithms and competitors. The regret is averaged over all users and over all runs. Under all graph models, **GraphUCB** outperforms its competitors consistently with a large margin. Also **GraphUCB-Local** consistently outperform competitor algorithms, with however slightly degraded performance with respect to **GraphUCB**. This is due to the approximation introduced by Eq. 7. **Gob.Lin** is a close runner since it is also based on the Laplacian regularized estimator, but it performs worse than the proposed algorithms, as already explained in the previous section. **CLUB** performs relative poor since there is no clear clusters in the graph. Nevertheless, it still outperforms **LinUCB** by grouping users into clusters in the early stage which speeds up the learning process. It is worth noting that the two subfigures depict the same algorithm for two different graph models (RBF and ER) with the same level of smoothness and sparsity. The trend of the cumulative regret is the same, highlighting that the algorithm is not affected by the graph model. This behavior is reinforced in Appendix I where we provide further results.

We are now interested in evaluating the performance of the proposed algorithms against different graph topologies, by varying signal smoothness and sparsity of graph (edge density) as follows:

Smoothness $[\gamma]$: We first generate a *RBF* graph. To control the smoothness, we vary $\gamma \in [0, 10]$.

Sparsity $[s]$: We first generate a *RBF* graph, then generate a smooth Θ via Eq. 19. To control the sparsity, we set a threshold $s \in [0, 1]$ on edge weights W_{ij} such that W_{ij} less than s are removed.

Figure 3 depicts the cumulative regret for different level of smoothness and sparsity of \mathcal{G} . **GraphUCB** and **GraphUCB-Local** show similar patterns (with **GraphUCB-Local** leading to higher regret due to the already commented approximation): (i) the smoother Θ the lower the regret, which is consistent

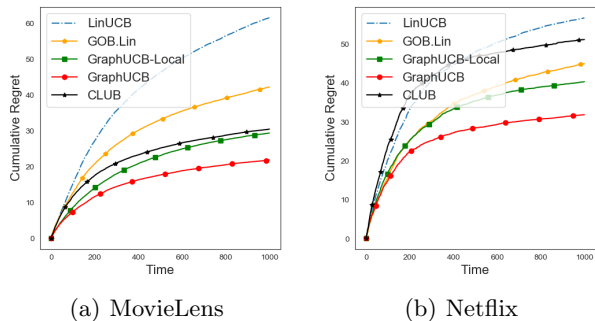


Figure 4: Performance on Real-World data.

with the Laplacian-regularized estimator of Eq. 4; (ii) denser graphs lead to lower regret since more connectivity provides more graph information which speeds up the learning process.

7.2 Experiments on Real-World Data

We then carry out experiments on two real-world datasets : **MovieLens** [Lam and Herlocker, 2006] and **Netflix** [Bennett et al., 2007]. We follow the data pre-processing steps in [Valko et al., 2014], described in details in Appendix K. The cumulative regret over time is depicted in Figure 4 for both datasets. Both **GraphUCB** and **GraphUCB-Local** outperform baseline algorithms in all cases. Similarly to the synthetic experiments, **LinUCB** performs poorly, while **Gob.Lin** shows a regret behavior more similar to the proposed algorithms. In the case of **MovieLens**, **CLUB** outperforms **Gob.Lin**. A close inspection of the data reveals that ratings provided by all users are highly concentrated. It means most users like a few set of movies. This is a good model for the clustering algorithm implemented in **CLUB**, hence the gain.

8 Conclusion

In this work, we propose a graph-based bandit algorithm **GraphUCB** and its scalable version **GraphUCB-Local**, both of which outperform the state-of-art bandit algorithms in terms of cumulative regret. On the theoretical side, we introduce a novel UCB embedding the graph structure in a natural way and show clearly that exploring the graph prior could reduce the cumulative regret. We demonstrate that the graph structure helps reduce the size of confidence set of the estimation of user features and the uncertainty of predicted payoff. As for future research directions, one possibility is to relax the assumption that the user graph is available and infer the graph from the data, ideally in a dynamic fashion.

Acknowledge

This work has been partially funded by Cisco Research Center.

References

- [Abbasi-Yadkori et al., 2011] Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320.
- [Agrawal and Goyal, 2013] Agrawal, S. and Goyal, N. (2013). Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135.
- [Alvarez et al., 2012] Alvarez, M. A., Rosasco, L., Lawrence, N. D., et al. (2012). Kernels for vector-valued functions: A review. *Foundations and Trends® in Machine Learning*, 4(3):195–266.
- [Auer, 2002] Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422.
- [Bennett et al., 2007] Bennett, J., Lanning, S., et al. (2007). The netflix prize. In *Proceedings of KDD cup and workshop*, volume 2007, page 35. New York, NY, USA.
- [Cesa-Bianchi et al., 2013] Cesa-Bianchi, N., Gentile, C., and Zappella, G. (2013). A gang of bandits. In *Advances in Neural Information Processing Systems*, pages 737–745.
- [Chapelle and Li, 2011] Chapelle, O. and Li, L. (2011). An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257.
- [Dani et al., 2008] Dani, V., Hayes, T. P., and Kakade, S. M. (2008). Stochastic linear optimization under bandit feedback.
- [Gentile et al., 2014] Gentile, C., Li, S., and Zappella, G. (2014). Online clustering of bandits. In *International Conference on Machine Learning*, pages 757–765.
- [Korda et al., 2016] Korda, N., Szörényi, B., and Shuai, L. (2016). Distributed clustering of linear bandits in peer to peer networks. In *Journal of machine learning research workshop and conference proceedings*, volume 48, pages 1301–1309. International Machine Learning Societ.
- [Lam and Herlocker, 2006] Lam, S. and Herlocker, J. (2006). Movielens data sets. *Department of Computer Science and Engineering at the University of Minnesota*.
- [Lattimore and Szepesvari, 2016] Lattimore, T. and Szepesvari, C. (2016). The end of optimism? an asymptotic analysis of finite-armed linear bandits. *arXiv preprint arXiv:1610.04491*.
- [Lattimore and Szepesvári, 2018] Lattimore, T. and Szepesvári, C. (2018). Bandit algorithms. *preprint*.
- [Li et al., 2010] Li, L., Chu, W., Langford, J., and Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670. ACM.
- [Li et al., 2016] Li, S., Karatzoglou, A., and Gentile, C. (2016). Collaborative filtering bandits. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 539–548. ACM.
- [Liu et al., 2018] Liu, B., Wei, Y., Zhang, Y., Yan, Z., and Yang, Q. (2018). Transferable contextual bandit for cross-domain recommendation. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [Nguyen and Lauw, 2014] Nguyen, T. T. and Lauw, H. W. (2014). Dynamic clustering of contextual multi-armed bandits. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 1959–1962. ACM.
- [Robbins, 1952] Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535.
- [Shuman et al., 2013] Shuman, D. I., Narang, S. K., Frossard, P., Ortega, A., and Vandergheynst, P. (2013). The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98.
- [Valko et al., 2014] Valko, M., Munos, R., Kveton, B., and Kocák, T. (2014). Spectral bandits for smooth graph functions. In *International Conference on Machine Learning*, pages 46–54.
- [Vaswani et al., 2017] Vaswani, S., Schmidt, M., and Lakshmanan, L. V. (2017). Horde of bandits using gaussian markov random fields. *arXiv preprint arXiv:1703.02626*.

- [Wu et al., 2016] Wu, Q., Wang, H., Gu, Q., and Wang, H. (2016). Contextual bandits in a collaborative environment. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 529–538. ACM.
- [Yang and Toni, 2018] Yang, K. and Toni, L. (2018). Graph-based recommendation system. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 798–802. IEEE.
- [Yankelevsky and Elad, 2016] Yankelevsky, Y. and Elad, M. (2016). Dual graph regularized dictionary learning. *IEEE Transactions on Signal and Information Processing over Networks*, 2(4):611–624.

Appendix A

Proof of Lemma 1

Proof.

$$\text{vec}(\hat{\Theta}_t) = (\mathbf{A}_t + \alpha L)^{-1} \Phi_t \mathbf{Y}_t \quad (20)$$

Let $\mathbf{M}_t = \mathbf{A}_t + \alpha L$ and $\mathbf{B}_t = \Phi_t \mathbf{Y}_t$, then

$$\text{vec}(\hat{\Theta}_t) = \mathbf{M}_t^{-1} \mathbf{B}_t \quad (21)$$

Express \mathbf{M}_t and \mathbf{B}_t in partitioned form. For instance, if $n = 2$

$$\mathbf{M}_t^{-1} = \begin{bmatrix} \mathbf{M}_{11,t}^{-1} & \mathbf{M}_{12,t}^{-1} \\ \mathbf{M}_{21,t}^{-1} & \mathbf{M}_{22,t}^{-1} \end{bmatrix} \quad \mathbf{B}_t = \begin{bmatrix} \mathbf{B}_{1,t} \\ \mathbf{B}_{2,t} \end{bmatrix} \quad (22)$$

Then

$$\hat{\theta}_{i,t} = \sum_{j=1}^n \mathbf{M}_{ij,t}^{-1} \mathbf{B}_{j,t} \quad (23)$$

In general case, $n \geq 2$,

$$\hat{\theta}_{i,t} = \sum_{j=1}^n \mathbf{M}_{ij,t}^{-1} \mathbf{B}_{j,t} \quad (24)$$

To obtain the expression of $\hat{\theta}_{i,t}$, we need the close form of $\mathbf{M}_{ij,t}^{-1}$. Given $\mathbf{M}_t = \mathbf{A}_t + \alpha L$ we have

$$\begin{aligned} \mathbf{M}_t^{-1} &= (\mathbf{A}_t + \alpha L)^{-1} \\ &= (\mathbf{A}_t \mathbf{A}_t^{-1} \mathbf{A}_t + \alpha L \mathbf{A}_t^{-1} \mathbf{A}_t)^{-1} \\ &= (\mathbf{I} + \alpha L \mathbf{A}_t^{-1}) \mathbf{A}_t^{-1} \\ &= \mathbf{A}_t^{-1} (\mathbf{I} + \alpha L \mathbf{A}_t^{-1})^{-1} \end{aligned} \quad (25)$$

This can be rewritten using Taylor expansion

$$\begin{aligned} &\mathbf{A}_t^{-1} (\mathbf{I} + \alpha L \mathbf{A}_t^{-1})^{-1} \\ &= \mathbf{A}_t^{-1} \mathbf{I} - \alpha L \mathbf{A}_t^{-1} + (\alpha L \mathbf{A}_t^{-1})^2 - \dots \\ &\approx \mathbf{A}_t^{-1} - \alpha \mathbf{A}_t^{-1} L \mathbf{A}_t^{-1} \end{aligned} \quad (26)$$

The last step keeps the first two terms only. So,

$$\mathbf{M}_t^{-1} \approx \mathbf{A}_t^{-1} - \alpha \mathbf{A}_t^{-1} L \mathbf{A}_t^{-1} \quad (27)$$

To obtain the expression of $\mathbf{M}_{ij,t}^{-1}$, we need the partitioned form of \mathbf{A}_t^{-1} and L .

Since \mathbf{A}_t is a block diagonal matrix, the inversion of it is the inversion of its block matrix. For instance, if $n = 2$

$$\mathbf{A}_t^{-1} = \begin{bmatrix} \mathbf{A}_{1,t}^{-1} & 0 \\ 0 & \mathbf{A}_{2,t}^{-1} \end{bmatrix} \quad (28)$$

and

$$L = \begin{bmatrix} \mathcal{L}_{11} \mathbf{I} & \mathcal{L}_{12} \mathbf{I} \\ \mathcal{L}_{21} \mathbf{I} & \mathcal{L}_{22} \mathbf{I} \end{bmatrix} \quad (29)$$

Given Eq. 27, in general case, $n \geq 2$, it is trivial to show

$$\mathbf{M}_{ij,t}^{-1} \approx \begin{cases} \mathbf{A}_{i,t}^{-1} - \alpha \mathbf{A}_{i,t}^{-1} \mathcal{L}_{ii} \mathbf{A}_{i,t}^{-1} & \text{when } i = j \\ -\alpha \mathbf{A}_{i,t}^{-1} \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} & \text{when } i \neq j \end{cases} \quad (30)$$

Finally, we have

$$\begin{aligned} \hat{\theta}_{i,t} &= \sum_{j=1}^n \mathbf{M}_{ij,t}^{-1} \mathbf{B}_{j,t} \\ &\approx (\mathbf{A}_{i,t}^{-1} - \alpha \mathbf{A}_{i,t}^{-1} \mathcal{L}_{ii} \mathbf{A}_{i,t}^{-1}) \mathbf{B}_{i,t} \\ &\quad - \alpha \mathbf{A}_{i,t}^{-1} \sum_{j \neq i} \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \mathbf{B}_{j,t} \\ &= \mathbf{A}_{i,t}^{-1} \mathbf{B}_{i,t} - \alpha \mathbf{A}_{i,t}^{-1} \sum_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \mathbf{B}_{j,t} \end{aligned} \quad (31)$$

We claim the approximation introduced in Eq. 31 is small. To see this, tracing back to Eq. 26 where the higher order terms are dropped. These higher order terms are negligible when t is large. This is because the Gram matrix \mathbf{A}_t grows with larger t , in turn \mathbf{A}_t^{-1} decreases. In such case, higher order terms would be incomparable with lower order terms.

To support the tightness of this approximation, we provide an empirical evidence in Figure 5. Red curve represents the single-user estimation error of The Laplacian-regularised estimator Eq. 4, while green curve is the estimation error of Lemma 1. Blue curve represents the estimation error of ridge regression, which is included as a benchmark. Clearly, Lemma 1 is a tight approximation of Eq. 4 and both converge faster than ridge regression due to the smoothness prior. \square




Figure 5: Approximation accuracy

Appendix B

Proof of Eq. 9

Proof. Let $\mathbf{A}_t = \Phi_t \Phi_t^T$ and $L = L \mathbf{I}$ and $\mathbf{M}_t = \mathbf{A}_t + \alpha L$. $\epsilon_t = [\eta_1, \eta_2, \dots, \eta_t]$

$$\text{vec}(\hat{\Theta}_t) = \mathbf{M}_t^{-1} \Phi_t^T \mathbf{Y}_t \quad (32)$$

The variance Σ_t is

$$\begin{aligned} \Sigma_t &= \text{Cov}(\text{vec}(\hat{\Theta}_t)) = \text{Cov}(\mathbf{M}_t^{-1} \Phi_t^T \mathbf{Y}_t) \\ &= \mathbf{M}_t^{-1} \Phi_t \text{Cov}(\mathbf{Y}_t) \Phi_t^T \mathbf{M}_t^{-1} \quad (33) \\ &= \sigma^2 \mathbf{M}_t^{-1} \mathbf{A}_t \mathbf{M}_t^{-1} \end{aligned}$$

where we use $\text{Cov}(\mathbf{Y}_t) = \sigma^2 \mathbf{I}$ since noise follows $\mathcal{N}(0, \sigma^2)$. Therefore, the precision matrix

$$\Lambda_t = \Sigma_t^{-1} = \frac{1}{\sigma^2} \mathbf{M}_t \mathbf{A}_t^{-1} \mathbf{M}_t \quad (34)$$

For simplicity, we assume $\sigma = 1$

$$\Lambda_t = \Sigma_t^{-1} = \mathbf{M}_t \mathbf{A}_t^{-1} \mathbf{M}_t \quad (35)$$

□

Appendix C

Proof of Eq. 10

Proof. Recall

$$\Lambda_t = \mathbf{M}_t \mathbf{A}_t^{-1} \mathbf{M}_t \quad (36)$$

and $\Lambda_{i,t} \in \mathbb{R}^{d \times d}$ is the i -th block matrix along the diagonal of Λ_t .

To get the expression of $\Lambda_{i,t}$, we need to express \mathbf{M}_t , \mathbf{A}_t^{-1} in partitioned form. For instance, $n = 2$

$$\mathbf{M}_t = \begin{bmatrix} \mathbf{M}_{11,t} & \mathbf{M}_{12,t} \\ \mathbf{M}_{21,t} & \mathbf{M}_{22,t} \end{bmatrix}, \quad \mathbf{A}_t^{-1} = \begin{bmatrix} \mathbf{A}_{1,t}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{2,t}^{-1} \end{bmatrix} \quad (37)$$

In general case $n \geq 2$, it is straightforward to see

$$\begin{aligned} \Lambda_{i,t} &= \prod_{j=1}^n \mathbf{M}_{ij,t} \mathbf{A}_{i,t}^{-1} \mathbf{M}_{ij,t} \\ &= \mathbf{M}_{ii,t} \mathbf{A}_{i,t}^{-1} \mathbf{M}_{ii,t} + \prod_{j \neq i} \mathbf{M}_{ij,t} \mathbf{A}_{j,t}^{-1} \mathbf{M}_{ji,t} \end{aligned} \quad (38)$$

From $\mathbf{M}_t = \mathbf{A}_t + \alpha L$, we know

$$\mathbf{M}_{ii,t} = \mathbf{A}_{i,t} + \alpha \mathcal{L}_{ii} \mathbf{I} \quad (39)$$

$$\mathbf{M}_{ij,t} = \alpha \mathcal{L}_{ij} \mathbf{I} \quad (40)$$

Hence,

$$\Lambda_{i,t} = \mathbf{A}_{i,t} + 2\alpha \mathcal{L}_{ii} \mathbf{I} + \alpha^2 \prod_{j=1}^n \mathcal{L}_{ij}^2 \mathbf{A}_{j,t}^{-1} \quad (41)$$

□

Appendix D

Proof of Lemma 2

Proof.

$$\mathcal{C}_t = \{\theta_i : \|\hat{\theta}_{i,t} - \theta_i\|_{i,t} \leq \beta_{i,t}\} \quad (42)$$

From Lemma 1, we have

$$\hat{\theta}_{i,t} \approx \mathbf{A}_{i,t}^{-1} \mathbf{X}_{i,t} \mathbf{Y}_{i,t} - \alpha \mathbf{A}_{i,t}^{-1} \prod_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \mathbf{X}_{j,t} \mathbf{Y}_{j,t} \quad (43)$$

Note that $\mathbf{Y}_{i,t} = \mathbf{X}_{i,t}^T \theta_i + \epsilon_{i,t}$ and $\mathbf{Y}_{j,t} = \mathbf{X}_{j,t}^T \theta_j + \epsilon_{j,t}$ where $\epsilon_{i,t} = [\eta_{i,1}, \dots, \eta_{i,T_{i,t}}]$ is the collection of noise associated with user i . $\mathcal{T}_{i,t}$ the set of time user i is selected during the time period from 1 to t .

Then we have

$$\begin{aligned} \hat{\theta}_{i,t} &= \mathbf{A}_{i,t}^{-1} \mathbf{X}_{i,t} (\mathbf{X}_{i,t}^T \theta_i + \epsilon_{i,t}) \\ &\quad - \alpha \mathbf{A}_{i,t}^{-1} \prod_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \mathbf{X}_{j,t} (\mathbf{X}_{j,t}^T \theta_j + \epsilon_{j,t}) \\ &= \theta_i + \mathbf{A}_{i,t}^{-1} \mathbf{X}_{i,t} \epsilon_{i,t} - \alpha \mathbf{A}_{i,t}^{-1} \prod_{j=1}^n \mathcal{L}_{ij} \theta_j \\ &\quad - \alpha \mathbf{A}_{i,t}^{-1} \prod_{j \neq i} \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \mathbf{X}_{j,t} \epsilon_{j,t} \end{aligned} \quad (44)$$

Hence

$$\begin{aligned} \hat{\theta}_{i,t} - \theta_i &= \mathbf{A}_{i,t}^{-1} \mathbf{X}_{i,t} \epsilon_{i,t} - \alpha \mathbf{A}_{i,t}^{-1} \prod_{j=1}^n \mathcal{L}_{ij} \theta_j \\ &\quad - \alpha \mathbf{A}_{i,t}^{-1} \prod_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \mathbf{X}_{j,t} \epsilon_{j,t} \end{aligned} \quad (45)$$

Denote $\xi_{i,t} = \mathbf{X}_{i,t} \epsilon_{i,t}$ and $\mathbf{V}_{i,t} = \mathbf{A}_{i,t} + \alpha \mathcal{L}_{ii} \mathbf{I}$, then,

$$\begin{aligned} \|\hat{\theta}_{i,t} - \theta_i\|_{i,t} &\leq \|\alpha \mathbf{A}_{i,t}^{-1} \prod_{j=1}^n \mathcal{L}_{ij} \theta_j\|_{i,t} \\ &\quad + \|\mathbf{A}_{i,t}^{-1} \xi_{i,t} - \alpha \mathbf{A}_{i,t}^{-1} \prod_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \xi_{j,t}\|_{i,t} \end{aligned} \quad (46)$$

$$\begin{aligned} &= \alpha \left\| \prod_{j=1}^n \mathcal{L}_{ij} \theta_j \right\|_{\mathbf{A}_{i,t}^{-1} \mathbf{V}_{i,t} \mathbf{A}_{i,t}^{-1}} \\ &\quad + \|\xi_{i,t} - \alpha \prod_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \xi_{j,t}\|_{\mathbf{A}_{i,t}^{-1} \mathbf{V}_{i,t} \mathbf{A}_{i,t}^{-1}} \end{aligned} \quad (47)$$

Here we apply $\|\mathbf{A}_{i,t}^{-1}(\cdot)\|_{i,t} = \|\cdot\|_{\mathbf{A}_{i,t}^{-1} \mathbf{V}_{i,t} \mathbf{A}_{i,t}^{-1}}$.

$$\leq \alpha \left\| \prod_{j=1}^n \mathcal{L}_{ij} \theta_j \right\|_{\mathbf{A}_{i,t}^{-1}} + \|\xi_{i,t} - \alpha \prod_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \xi_{j,t}\|_{\mathbf{A}_{i,t}^{-1}} \quad (48)$$

Here we use $\|\cdot\|_{\mathbf{A}_{i,t}^{-1}} \leq \|\cdot\|_{\mathbf{A}_{i,t}^{-1}}$.

$$\leq \alpha \sum_{j=1}^n \mathcal{L}_{ij} \theta_j \| \mathbf{v}_{i,t}^{-1} \| + \| \boldsymbol{\xi}_{i,t} - \alpha \sum_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \boldsymbol{\xi}_{i,t} \|_{\mathbf{v}_{i,t}^{-1}} \quad (49)$$

Here we use $\|\cdot\|_{\mathbf{A}_{i,t}^{-1}} \leq \|\cdot\|_{\mathbf{V}_{i,t}^{-1}}$.

$$\leq \alpha \sum_{j=1}^n \mathcal{L}_{ij} \theta_j \| \mathbf{v}_{i,t}^{-1} \| + \| \boldsymbol{\xi}_{i,t} \|_{\mathbf{V}_{i,t}^{-1}} \quad (50)$$

Here we use

$$\| \boldsymbol{\xi}_{i,t} - \alpha \sum_{j=1}^n \mathcal{L}_{ij} \mathbf{A}_{j,t}^{-1} \boldsymbol{\xi}_{i,t} \|_{\mathbf{V}_{i,t}^{-1}} \leq \| \boldsymbol{\xi}_{i,t} \|_{\mathbf{V}_{i,t}^{-1}} \quad (51)$$

To support Eq. 51, we provide an empirical evidence in Figure 6.

Denote $\boldsymbol{\Delta}_i = \sum_{j=1}^n \mathcal{L}_{ij} \theta_j$, Finally, we have

$$\| \hat{\boldsymbol{\theta}}_i - \boldsymbol{\theta}_i \|_{i,t} \leq \alpha \| \boldsymbol{\Delta}_i \|_{\mathbf{V}_{i,t}^{-1}} + \| \boldsymbol{\xi}_{i,t} \|_{\mathbf{V}_{i,t}^{-1}} \quad (52)$$

According to Theorem 2 in [Abbasi-Yadkori et al., 2011], we have the following upper bound holds with probability $1 - \delta$ for $\delta \in [0, 1]$.

$$\| \boldsymbol{\xi}_{i,t} \|_{\mathbf{V}_{i,t}^{-1}} \leq \sigma \sqrt{2 \log \frac{|\mathbf{V}_{i,t}|^{1/2}}{\delta |\alpha \mathbf{I}|^{1/2}}} \quad (53)$$

In addition, we know

$$\alpha \| \boldsymbol{\Delta}_i \|_{\mathbf{V}_{i,t}^{-1}} \leq \sqrt{\alpha} \| \boldsymbol{\Delta}_i \|_2 \quad (54)$$

where we use $\| \boldsymbol{\Delta}_i \|_{\mathbf{V}_{i,t}^{-1}}^2 \leq \frac{1}{\lambda_{\min}} (\mathbf{V}_{i,t}) \| \boldsymbol{\Delta}_i \|_2^2 \leq \frac{1}{\alpha} \| \boldsymbol{\Delta}_i \|_2$, which means $\alpha \| \boldsymbol{\Delta}_i \|_{\mathbf{V}_{i,t}^{-1}} \leq \sqrt{\frac{\alpha}{\alpha}} \| \boldsymbol{\Delta}_i \|_2 = \sqrt{\alpha} \| \boldsymbol{\Delta}_i \|_2$.

Finally, combine the above two upper bounds, we have the following upper bound holds with probability $1 - \delta$ for $\delta \in [0, 1]$.

$$\| \hat{\boldsymbol{\theta}}_i - \boldsymbol{\theta}_i \|_{i,t} \leq \sigma \sqrt{2 \log \frac{|\mathbf{V}_{i,t}|^{1/2}}{\delta |\alpha \mathbf{I}|^{1/2}}} + \sqrt{\alpha} \| \boldsymbol{\Delta}_i \|_2 \quad (55)$$

which means

$$\beta_{i,t} = \sigma \sqrt{2 \log \frac{|\mathbf{V}_{i,t}|^{1/2}}{\delta |\alpha \mathbf{I}|^{1/2}}} + \sqrt{\alpha} \| \boldsymbol{\Delta}_i \|_2 \quad (56)$$

□

The green curve represents the LHS term in Eq. 51, while the blue curve represents the RHS term. Clearly, the blue curve is above the green curve and they converges together with large t .




Figure 6: Noise term approximation

Appendix E

Pseudocode of GraphUCB-Local

Algorithm 2: GraphUCB-Local

Input : α, T, L, δ

Initialization : For any $i \in \{1, 2, \dots, n\}$

$\hat{\boldsymbol{\theta}}_{0,i} = \mathbf{0} \in \mathbb{R}^d, \boldsymbol{\Lambda}_{0,i} = \mathbf{0} \in \mathbb{R}^{d \times d},$

$\mathbf{A}_{0,i} = \mathbf{0} \in \mathbb{R}^{d \times d}, \beta_{i,t} = 0.$

for $t \in [1, T]$ **do**

 User index i is selected

 1. $\mathbf{A}_{i,t} \leftarrow \mathbf{A}_{i,t-1} + \mathbf{x}_t \mathbf{x}_t^T$.

 2. $\mathbf{A}_{j,t} \leftarrow \mathbf{A}_{j,t-1}, \forall j \neq i.$

 3. Update $\boldsymbol{\Lambda}_{i,t}$ via Eq. 10.

 4. Select \mathbf{x}_t via Eq. 12

 where $\beta_{i,t}$ is defined in Eq. 11.

 5. Receive the payoff y_t .

 6. Update $\hat{\boldsymbol{\theta}}_{i,t}$ via Lemma. 1 if $i = i_t$.

 7. $\hat{\boldsymbol{\theta}}_{j,t} \leftarrow \hat{\boldsymbol{\theta}}_{j,t-1} \quad \forall j \neq i_t.$

end

Appendix F

Proof of Lemma 3: Recall $\Psi_{i,T} = \frac{\sum_{t \in \mathcal{T}_{i,T}} \mathbf{x}_t \mathbf{x}_t^T}{\sum_{t \in \mathcal{T}_{i,T}} \mathbf{x}_t \mathbf{x}_t^T} \mathbf{I}_{i,T}^{-1}$,

where $\mathcal{T}_{i,T}$ is the set of time user i is served up to time T , $\mathbf{A}_{i,\tau} = \sum_{t \in \mathcal{T}_{i,\tau}} \mathbf{x}_t \mathbf{x}_t^T$, $\mathbf{V}_{i,\tau} = \mathbf{A}_{i,\tau} + \alpha \mathcal{L}_{ii} \mathbf{I}$ and $\boldsymbol{\Lambda}_{i,\tau}$ defined in Eq. 10. Without loss of generality, assume $\| \mathbf{x}_\tau \|_2 \leq 1$ for any $\tau \leq T$.

Proof.

$$\sum_{\tau \in \mathcal{T}_{i,T}} \| \mathbf{x}_\tau \|_{\mathbf{V}_{i,\tau}^{-1}}^2 \leq (1 + \max_{\tau \in \mathcal{T}_{i,T}} \| \mathbf{x}_\tau \|_2) \log |\mathbf{V}_{i,\tau}| \quad (57)$$

In the same fashion

$$\sum_{\tau \in \mathcal{T}_{i,T}} \| \mathbf{x}_\tau \|_{\boldsymbol{\Lambda}_{i,\tau}^{-1}}^2 \leq (1 + \max_{\tau \in \mathcal{T}_{i,T}} \| \mathbf{x}_\tau \|_2) \log |\boldsymbol{\Lambda}_{i,\tau}| \quad (58)$$

Since we assume $\|\mathbf{x}_\tau\|_2 \leq 1$ for any $\tau \leq T$ and

$$\Psi_{i,T} = \frac{\prod_{\tau=2T_{i,T}}^T \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2}{\prod_{\tau=2T_{i,T}}^T \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2} \quad (59)$$

Given

$$\begin{aligned} \mathbf{V}_{i,\tau} &= \mathbf{A}_{i,\tau} + \alpha \mathbf{I} \\ \mathbf{A}_{i,\tau} &= \mathbf{A}_{i,\tau} + 2\alpha \mathcal{L}_{ii} \mathbf{I} + \alpha^2 \prod_{j=1}^n \mathcal{L}_{ij}^2 \mathbf{A}_{j,\tau} \end{aligned} \quad (60)$$

Assume⁵ $\mathbf{A}_{i,\tau}$ (and $\mathbf{A}_{j,\tau}$) are positive semi-definite for all τ . we know that $\mathbf{A}_{i,\tau} > \mathbf{V}_{i,\tau}$, therefore

$$\mathbf{A}_{i,\tau}^{-1} < \mathbf{V}_{i,\tau}^{-1} \quad (61)$$

holds for any $\tau \leq T$, which means

$$\|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}} < \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}} \quad (62)$$

holds for any $\tau \leq T$.

Thus, we have

$$\prod_{\tau=2T_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2 < \prod_{\tau=2T_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2 \quad (63)$$

this means

$$\Psi_{i,T} = \frac{\prod_{\tau=2T_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2}{\prod_{\tau=2T_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2} < 1 \quad (64)$$

In addition, since $\|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2 > 0$ and $\|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2 > 0$,

This must hold

$$\Psi_{i,T} > 0 \quad (65)$$

Combine all together, we have

$$\Psi_{i,T} \in (0, 1) \quad (66)$$

Furthermore, as $\mathbf{A}_{j,\tau}^{-1}$ decreases over time τ , $\mathbf{A}_{i,\tau} \rightarrow \mathbf{A}_{i,T}$. It means $\Psi_{i,T} \rightarrow 1$. \square

Appendix G

Proof of Theorem 1

Proof. First, we show the instantaneous regret at time t can be upper bounded by $2\beta_{i,t} \|\mathbf{x}_{i,t}\|_{\mathbf{V}_{i,t}^{-1}}$ where \mathbf{x}_t is

⁴For isolated node, we set $L_{ii} = 1$, $L_{ij} = 0; j \notin i$.

⁵This can be ensured trivially by adding \mathbf{I} to $\mathbf{A}_{i,\tau}$ with a small ϵ .

the arm selected by the learner at time t for user i . $\mathbf{x}_{i,t}$ is the optimal arm for user i .

$$\begin{aligned} r_{i,t} &= \mathbf{x}_{i,t}^T \boldsymbol{\theta}_i - \mathbf{x}_t^T \boldsymbol{\theta}_i \\ &\leq \mathbf{x}_t^T \hat{\boldsymbol{\theta}}_{i,t} + \beta_{i,t} \|\mathbf{x}_t\|_{\mathbf{V}_{i,t}^{-1}} - \mathbf{x}_{i,t}^T \boldsymbol{\theta}_i \\ &\leq \mathbf{x}_t^T \hat{\boldsymbol{\theta}}_{i,t} + \beta_{i,t} \|\mathbf{x}_t\|_{\mathbf{V}_{i,t}^{-1}} - \mathbf{x}_t^T \hat{\boldsymbol{\theta}}_{i,t} + \beta_{i,t} \|\mathbf{x}_t\|_{\mathbf{V}_{i,t}^{-1}} \\ &= 2\beta_{i,t} \|\mathbf{x}_t\|_{\mathbf{V}_{i,t}^{-1}} \end{aligned} \quad (67)$$

where we use the principle of optimistic

$$\mathbf{x}_{i,t}^T \boldsymbol{\theta}_i + \beta_{i,t} \|\mathbf{x}_{i,t}\|_{\mathbf{V}_{i,t}^{-1}} \leq \mathbf{x}_t^T \hat{\boldsymbol{\theta}}_{i,t} + \beta_{i,t} \|\mathbf{x}_t\|_{\mathbf{V}_{i,t}^{-1}} \quad (68)$$

and

$$\mathbf{x}_t^T \hat{\boldsymbol{\theta}}_{i,t} \leq \mathbf{x}_t^T \boldsymbol{\theta}_i + \beta_{i,t} \|\mathbf{x}_t\|_{\mathbf{V}_{i,t}^{-1}} \quad (69)$$

Next, we drive a upper bound of the cumulative regret of user i up to T

$$\begin{aligned} R_{i,T} &= \sum_{\tau=2T_{i,T}}^T r_{i,\tau} \leq \frac{|\mathcal{T}_{i,T}|}{\tau=2T_{i,T}} \times \frac{r_{i,\tau}^2}{r_{i,\tau}^2} \\ &\leq \frac{|\mathcal{T}_{i,T}|}{\tau=2T_{i,T}} \times \frac{4\beta_{i,T}^2 \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2}{\tau=2T_{i,T}} \\ &\leq 2\beta_{i,T} \frac{|\mathcal{T}_{i,T}|}{\tau=2T_{i,T}} \times \frac{\|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2}{\tau=2T_{i,T}} \\ &\leq 2\beta_{i,T} \frac{|\mathcal{T}_{i,T}|}{\tau=2T_{i,T}} \times \frac{\min(1, \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2)}{\tau=2T_{i,T}} \end{aligned} \quad (70)$$

where we use $\beta_{i,T} \geq \beta_{i,\tau}$ since $\beta_{i,\tau}$ is an increasing function over τ and $r_{i,\tau} \leq 2$ since we assume payoff $\mathbf{x}^T \boldsymbol{\theta}_i \in [-1, 1]$.

According to Lemma 11 in [Abbasi-Yadkori et al., 2011], we have

$$\begin{aligned} \prod_{\tau=2T_{i,T}} \min(1, \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2) &\leq 2 \log \frac{|\mathbf{V}_{i,T}|}{|\alpha \mathbf{I}|} \\ &\leq 2 \frac{d |\mathcal{T}_{i,T}| \log(\alpha + |\mathcal{T}_{i,T}|/d)}{d |\mathcal{T}_{i,T}| \log(\alpha + |\mathcal{T}_{i,T}|/d)} \end{aligned} \quad (71)$$

where $\mathbf{A}_{i,T} = \prod_{\tau=2T_{i,T}} \mathbf{x}_\tau \mathbf{x}_\tau^T$, $\mathbf{V}_{i,T} = \mathbf{A}_{i,T} + \alpha \mathcal{L}_{ii} \mathbf{I}$ and $\mathcal{L}_{ii} = 1$.

Recall in Lemma 3, we define

$$\Psi_{i,T} = \frac{\prod_{\tau=2T_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2}{\prod_{\tau=2T_{i,T}} \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2} \quad (72)$$

Therefore

$$\begin{aligned} \prod_{\tau=2T_{i,T}} \min(1, \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}^{-1}}^2) &= \Psi_{i,T} \prod_{\tau=2T_{i,T}} \min(1, \|\mathbf{x}_\tau\|_{\mathbf{V}_{i,\tau}}^2) \\ &\leq 2\Psi_{i,T} \frac{d |\mathcal{T}_{i,T}| \log(\alpha + |\mathcal{T}_{i,T}|/d)}{d |\mathcal{T}_{i,T}| \log(\alpha + |\mathcal{T}_{i,T}|/d)} \end{aligned} \quad (73)$$

From Eq. 56, we have

$$\beta_{i,T} = \sigma \sqrt{\frac{2 \log \frac{|\mathbf{V}_{i,T}|^{1/2}}{\delta |\alpha \mathbf{I}|^{1/2}} + \sqrt{\alpha} \|\Delta_i\|_2}{\delta}} \quad (74)$$

According to Theorem 2 in [Abbasi-Yadkori et al., 2011],

$$\sigma \sqrt{\frac{2 \log \frac{|\mathbf{V}_{i,T}|^{1/2}}{\delta |\alpha \mathbf{I}|^{1/2}}}{\delta}} \leq \sigma \sqrt{\frac{d \log \frac{1 + |\mathcal{T}_{i,T}|/\alpha}{\delta}}{\delta}} \leq \mathcal{O}(\sqrt{d \log |\mathcal{T}_{i,T}|}) \quad (75)$$

Hence,

$$\beta_{i,t} \leq \mathcal{O}(\sqrt{d \log |\mathcal{T}_{i,T}|} + \sqrt{\alpha} \|\Delta_i\|_2) \quad (76)$$

Combine this with Eq. 73 and Eq. 70, we have

$$\begin{aligned} R_{i,T} &\leq \mathcal{O}(\sqrt{d \log |\mathcal{T}_{i,T}|} + \sqrt{\alpha} \|\Delta_i\|_2) \times \\ &\quad \mathcal{O}(\sqrt{d |\mathcal{T}_{i,T}| \log(|\mathcal{T}_{i,T}|)}) \\ &\leq \tilde{\mathcal{O}}(\Psi_{i,T} \sqrt{d |\mathcal{T}_{i,T}|}) \end{aligned} \quad (77)$$

where the constant term $\sqrt{\alpha} \|\Delta_i\|_2$ and logarithmic terms are hidden.

Assume users are served uniformly, i.e., $\mathcal{T}_{i,T} = T/n$. Then, over the time horizon T , the total cumulative regret experienced by all users satisfies the following upper bound with probability $1 - \delta$ with $\delta \in [0, 1]$.

$$\begin{aligned} R_T &= \sum_{i=1}^{\mathcal{X}} R_{i,T} = \sum_{i=1}^{\mathcal{X}} \tilde{\mathcal{O}}(\Psi_{i,T} \sqrt{d |\mathcal{T}_{i,T}|}) \\ &= \sum_{i=1}^{\mathcal{X}} \tilde{\mathcal{O}}(\Psi_{i,T} \sqrt{d T/n}) \\ &\leq \tilde{\mathcal{O}}(nd \sqrt{T/n} \max_{i \in \mathcal{U}} \Psi_{i,T}) \\ &= \tilde{\mathcal{O}}(d\sqrt{Tn} \max_{i \in \mathcal{U}} \Psi_{i,T}) \end{aligned} \quad (78)$$

□

Appendix H

Given $L = \mathbf{D}^{-1} \mathbf{L}$ is the random-walk graph Laplacian, where $\mathbf{L} = \mathbf{D} - \mathbf{W}$ is the combinatorial Laplacian. Recall $\Theta = [\theta_1, \theta_2, \dots, \theta_n] \in \mathbb{R}^{n \times d}$ contains user features $\theta_i \in \mathbb{R}^d$ in rows. Then, the quadratic Laplacian form can be expressed in the following way:

$$\text{tr}(\Theta^T L \Theta) = \sum_{k=1}^{\mathcal{X}} \sum_{i,j} \frac{1}{4} \frac{W_{ij}}{D_{ii}} + \frac{W_{ji}}{D_{jj}} \|\theta_{ik} - \theta_{jk}\|^2 \quad (79)$$

Proof.

$$\text{tr}(\Theta^T L \Theta) = \sum_{k=1}^{\mathcal{X}} \theta_{::k}^T L \theta_{::k} \quad (80)$$

where $\theta_{::k} \in \mathbb{R}^n$ is the k -th column of Θ .

Note that $L = \mathbf{D}^{-1} \mathbf{L}$ is an asymmetric matrix with off-diagonal element $\mathcal{L}_{ij} = -\frac{W_{ij}}{D_{ii}}$ and $\mathcal{L}_{ji} = -\frac{W_{ji}}{D_{jj}}$ and on-diagonal element $\mathcal{L}_{ii} = 1$.

From elementary linear algebra, we know that

$$L = \frac{L + L^T}{2} + \frac{L - L^T}{2} \quad (81)$$

Then

$$\begin{aligned} \text{tr}(\Theta^T L \Theta) &= \text{tr}(\Theta^T (\frac{L + L^T}{2} + \frac{L - L^T}{2}) \Theta) \\ &= \text{tr}(\Theta^T (\frac{L + L^T}{2}) \Theta) + \text{tr}(\Theta^T (\frac{L - L^T}{2}) \Theta) \\ &= \text{tr}(\Theta^T (\frac{L + L^T}{2}) \Theta) \end{aligned} \quad (82)$$

where

$$\text{tr}(\Theta^T (\frac{L - L^T}{2}) \Theta) = 0 \quad (83)$$

To see this

$$\text{tr}(\Theta^T (\frac{L - L^T}{2}) \Theta) = \sum_{k=1}^{\mathcal{X}} \theta_{::k}^T (\frac{L - L^T}{2}) \theta_{::k} \quad (84)$$

for any $k \in \{1, 2, \dots, d\}$

$$\theta_{::k}^T (\frac{L - L^T}{2}) \theta_{::k} = 0 \quad (85)$$

To see this, assume $p = \theta_{::k}^T (\frac{L - L^T}{2}) \theta_{::k}$, then

$$\begin{aligned} p &= \theta_{::k}^T (\frac{L - L^T}{2}) \theta_{::k} \\ &= \theta_{::k}^T (\frac{L - L^T}{2}) \theta_{::k}^T \\ &= -p \end{aligned} \quad (86)$$

where $\frac{L - L^T}{2} = -\frac{L^T - L}{2}$. So $p = 0$.

Then,

$$\text{tr}(\Theta^T L \Theta) = \text{tr}(\Theta^T (\frac{L + L^T}{2}) \Theta) \quad (87)$$

where the off-diagonal element, $i \neq j$, of $\frac{L + L^T}{2}$ is $\frac{1}{2} \frac{W_{ij}}{D_{ii}} + \frac{W_{ji}}{D_{jj}}$ and the on-diagonal element is 1.

Therefore

$$\text{tr}(\Theta^T L \Theta) = \sum_{k=1}^{\mathcal{X}} \sum_{i,j} \frac{1}{4} \frac{W_{ij}}{D_{ii}} + \frac{W_{ji}}{D_{jj}} \|\theta_{ik} - \theta_{jk}\|^2 \quad (88)$$

□

Appendix I

To simulate \mathcal{G} , we follow two random graph models commonly used in the network science community: 1) Radial basis function (*RBF*) model, a weighted fully connected graph, with edge weights $W_{ij} = \exp(-\rho\|\theta_i - \theta_j\|^2)$; 2) Erdős Rényi (ER) model, an unweighted graph, in which each edge is generated independently and randomly with probability p . 3) *Barabasi-Albert* (BA) model, an unweighted graph initialized with a connected graph with m nodes. Then, a new node is added to the graph sequentially with m edges connected to existing nodes following the rule of preferential attachment where existing nodes with more edges has more probability to be connected by the new node; 4) *Watts-Strogatz* (WS) model, an unweighted graph, which is a m -regular graph with edges randomly rewired with probability p . For each graph model, different topologies can be generated, leading to different level of sparsity and smoothness as show in the following

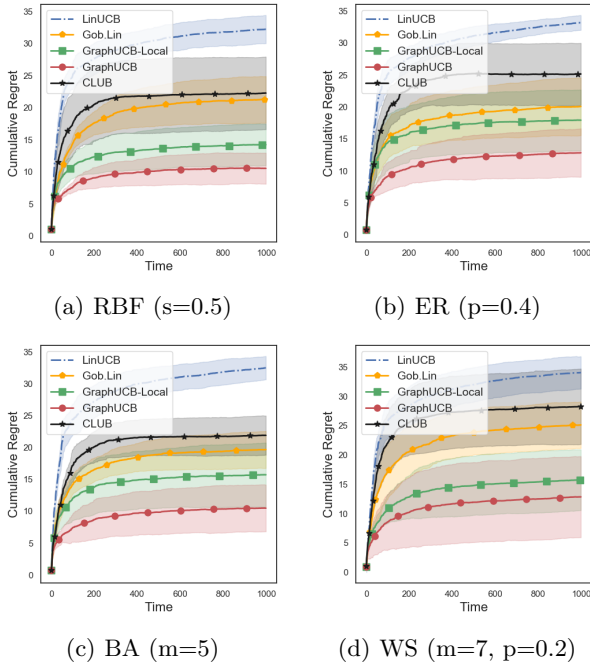



Figure 7: Performance with respect to Graph Types

Comparing sub-figures in Figure 7 shows that with the same level of sparsity and smoothness the effect of topology of graph performance seems to be unnoticeable. The generated graphs are shown in Figure 9. This is also confirmed in Figure 8, in which we generate graphs with different topology but with the same level connectivity. Algorithms are test on smoothness level.



(a) GraphUCB

(b) GraphUCB-Local

Figure 8: Performance of proposed algorithms on Graph models




Figure 9: Graphs

Appendix J

We test the performance of the proposed algorithms on the basis of graph properties such as smoothness, sparsity, p in ER graph, m in BA graph, m and p in WS graph.

Smoothness [γ]: We first generate a *RBF* graph. To control the smoothness, we vary $\gamma \in [0, 10]$. The term $sm = \text{tr}(\Theta^T L \Theta)$ measures the corresponding smoothness level. To ensure the comparison is fair, Θ is normalized to be $\|\Theta\|_2 = n$.

Sparsity [s]: We first generate a *RBF* graph, then generate a smooth Θ via Eq. 19. To control the sparsity, we set a threshold $s \in [0, 1]$ on edge weights W_{ij} such that W_{ij} less than s are removed. The term $sp = \frac{\text{number of edges}}{n(n-1)}$ measures the corresponding level

of sparsity, where $n(n-1)$ is the number of edges of a fully connected graph.

Results are shown in Figure 10. In sub-figure (a) and (b), graph-based algorithms show similar pattern. Smoother signal leads to less regret because of the Laplacian-regularizer estimator in Eq. 4. Sparse graph leads to more regret as less connectivity provides less graph information. This is also confirmed by sub-figure (c), p in ER controls the probability of edge. Small p leads to sparse graph, in turn more regret.

The dimension is set as $d = 10$. Next, we create the user graph \mathcal{G} from \mathbf{U} via *RBF* kernel. **Netflix** contains rating of 480k users on 18k movies. We process the dataset in the same way as **MovieLens**. In both datasets, original ratings range from 0 to 5, we normalize them into $[0, 1]$. After the data pre-processing, we sample 50 users and test algorithms over $T = 1000$.




Figure 11: Histogram of signals in MovieLens (a) and Netflix (b).

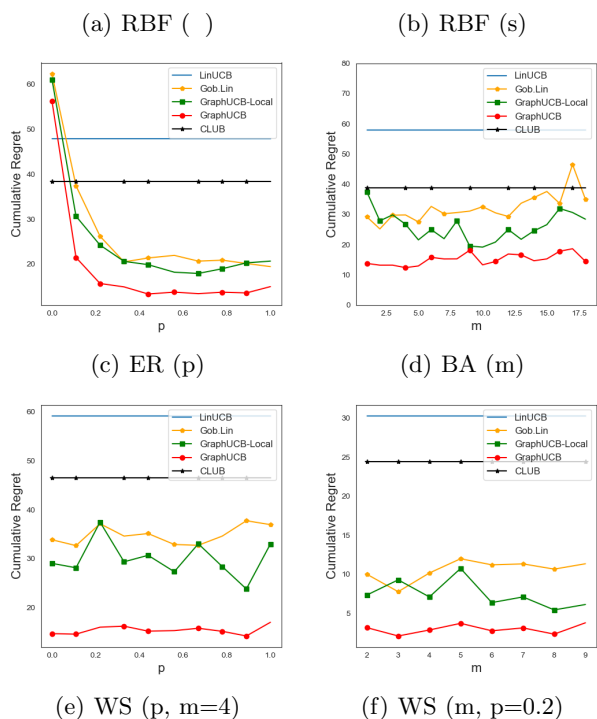


Figure 10: Performance on Graph Properties

Figure 11 shows the distribution of ratings in **MovieLens** and **Netflix**. Ratings in **MovieLens** are highly concentrated which means a large number of users like a few set of movies. They show similar performance.

Appendix K

MovieLens contains 6k users and their ratings on 40k movies. Since every user does not give ratings on all movies, there are a large amount of missing ratings. We factorize the rating matrix via $\mathbf{M} = \mathbf{U}\mathbf{X}$ to fill the missing values where \mathbf{U} contain users' latent vectors in rows and \mathbf{X} contain movies' latent features in columns.