Collaborative Filtering Bandits

Shuai Li University of Insubria, Italy shuaili.sli@gmail.com

Alexandros Karatzoglou Telefónica Research, Spain alexk@tid.es Claudio Gentile University of Insubria, Italy claudio.gentile@uninsubria.it

ABSTRACT

Classical collaborative filtering, and content-based filtering methods try to learn a static recommendation model given training data. These approaches are far from ideal in highly dynamic recommendation domains such as news recommendation and computational advertisement, where the set of items and users is very fluid. In this work, we investigate an adaptive clustering technique for content recommendation based on exploration-exploitation strategies in contextual multi-armed bandit settings. Our algorithm takes into account the collaborative effects that arise due to the interaction of the users with the items, by dynamically grouping users based on the items under consideration and, at the same time, grouping items based on the similarity of the clusterings induced over the users. The resulting algorithm thus takes advantage of preference patterns in the data in a way akin to collaborative filtering methods. We provide an empirical analysis on medium-size real-world datasets, showing scalability and increased prediction performance (as measured by click-through rate) over state-of-the-art methods for clustering bandits. We also provide a regret analysis within a standard linear stochastic noise setting.

Keywords

Filtering and Recommending; Recommender Systems; Online Learning; Collaborative Filtering; Clustering; Bandits; Regret; Computational Advertising

1. INTRODUCTION

Recommender Systems are an essential part of many successful on-line businesses, from e-commerce to on-line streaming, and beyond. Moreover, Computational Advertising can be seen as a recommendation problem where the user preferences highly depend on the current *context*. In fact, many recommendation domains such as Youtube video recommendation or news recommendation *do not* fit the classical description of a recommendation scenario, whereby a set of users with essentially fixed preferences interact with a fixed

SIGIR '16, July 17-21, 2016, Pisa, Italy

© 2016 ACM. ISBN 978-1-4503-4069-4/16/07...\$15.00

DOI: http://dx.doi.org/10.1145/2911451.2911548

set of items. In this classical setting, the well-known *cold-start* problem, namely, the lack of accumulated interactions by users on items, needs to be addressed, for instance, by turning to *hybrid* recommendation methods (e.g., [19]). In practice, many relevant recommendation domains are dynamic, in the sense that user preferences and the set of active users change with time. Recommendation domains can be distinguished by how much and how often user preferences and content universe change (e.g., [24]). In highly dynamic recommendation domains, such as news, ads and videos, active users and user preferences are fluid, hence classical collaborative filtering-type methods, such as Matrix or Tensor-Factorization break down. In these settings, it is essential for the recommendation method to adapt to the shifting preference patterns of the users.

Exploration-exploitation methods, a.k.a. multi-armed bandits, have been shown to be an excellent solution for these dynamic domains (see, e.g., the news recommendation evidence in [23]). While effective, standard contextual bandits do not take collaborative information into account, that is, users who have interacted with similar items in the past will not be deemed to have similar taste based on this fact alone, while items that have been chosen by the same group of users will also not be considered as similar. It is this significant limitation in the current bandit methodology that we try to address in this work. Past efforts on this problem were based on using online clustering-like algorithms on the graph or network structure of the data in conjunction with multi-armed bandit methods (see Section 3).

Commercial large scale search engines and information retrieval systems are examples of highly dynamic environments where users and items could be described in terms of their membership in some preference cluster. For instance, in a music recommendation scenario, we may have groups of listeners (the users) clustered around music genres, with the clustering changing across different genres. On the other hand, the individual songs (the items) could naturally be grouped by sub-genre or performer based on the fact that they tend to be preferred by the same group of users. Evidence has been collected which suggests that, at least in specific recommendation scenarios, like movie recommendation, data are well modeled by clustering at both user and item sides (e.g., [31]).

In this paper, we introduce a Collaborative Filtering based stochastic multi-armed bandit method that allows for a flexible and generic integration of information of users and items interaction data by alternatively clustering over both user and item sides. Specifically, we describe and analyze an

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

adaptive and efficient clustering of bandit algorithm that can perform collaborative filtering, named COFIBA (pronounced as "coffee bar"). Importantly enough, the clustering performed by our algorithm relies on sparse graph representations, avoiding expensive matrix factorization techniques. We adapt COFIBA to the standard setting of sequential content recommendation known as (contextual) multi-armed bandits (e.g., [5]) for solving the canonical exploration vs. exploitation dilemma.

Our algorithm works under the assumption that we have to serve content to users in such a way that each content *item* determines a clustering over users made up of relatively few groups (compared to the total number of users), within which users tend to react similarly when that item gets recommended. However, the clustering over users need not be the same across different items. Moreover, when the universe of items is large, we also assume that the items might be clustered as a function of the clustering they determine over users, in such a way that the number of *distinct* clusterings over users induced by the items is also relatively small compared to the total number of available items.

Our method aims to exploit collaborative effects in a bandit setting in a way akin to the way co-clustering techniques are used in batch collaborative filtering. Bandit methods also represent one of the most promising approaches to the research community of recommender systems, for instance in tackling the cold-start problem (e.g., [25]), whereby the lack of data on new users leads to suboptimal recommendations. An exploration approach in these cases seems very appropriate.

We demonstrate the efficacy of our dynamic clustering algorithm on three benchmark and real-world datasets. Our algorithm is scalable and exhibits significant increased prediction performance over the state-of-the-art of clustering bandits. We also provide a regret analysis of the \sqrt{T} -style holding with high probability in a standard stochastically linear noise setting.

2. LEARNING MODEL

We assume that the user behavior similarity is encoded by a family of clusterings depending on the specific feature (or context, or item) vector \boldsymbol{x} under consideration. Specifically, we let $\mathcal{U} = \{1, \dots, n\}$ represent the set of n users. Then, given $\boldsymbol{x} \in \mathbb{R}^{d}$, set \mathcal{U} can be partitioned into a small number $m(\boldsymbol{x})$ of clusters $U_1(\boldsymbol{x}), U_2(\boldsymbol{x}), \ldots, U_{m(\boldsymbol{x})}(\boldsymbol{x})$, where $m(\boldsymbol{x})$ is upper bounded by a constant m, independent of \boldsymbol{x} , with m being much smaller than n. (The assumption $m \ll n$ is not strictly required but it makes our algorithms more effective, and this is actually what we expect our datasets to comply with.) The clusters are such that users belonging to the same cluster $U_i(\boldsymbol{x})$ tend to have similar behavior w.r.t. feature vector \boldsymbol{x} (for instance, they both like or both dislike the item represented by x), while users lying in different clusters have significantly different behavior. The mapping $\boldsymbol{x} \to \{U_1(\boldsymbol{x}), U_2(\boldsymbol{x}), \dots, U_{m(\boldsymbol{x})}(\boldsymbol{x})\}$ specifying the actual partitioning of the set of users \mathcal{U} into the clusters determined by \boldsymbol{x} (including the number of clusters $m(\boldsymbol{x})$ and its upper bound m), as well as the common user behavior within each cluster are unknown to the learning system, and have to be inferred based on user feedback.

For the sake of simplicity, this paper takes the simple viewpoint that clustering over users is determined by linear functions $\boldsymbol{x} \to \boldsymbol{u}_i^\top \boldsymbol{x}$, each one parameterized by an un-

known vector $\boldsymbol{u}_i \in \mathbb{R}^d$ hosted at user $i \in \mathcal{U}$, in such a way that if users i and i' are in the same cluster w.r.t. \boldsymbol{x} then $\boldsymbol{u}_i^{\top} \boldsymbol{x} = \boldsymbol{u}_{i'}^{\top} \boldsymbol{x}$, while if i and i' are in different clusters w.r.t. \boldsymbol{x} then $|\boldsymbol{u}_i^{\top} \boldsymbol{x} - \boldsymbol{u}_{i'}^{\top} \boldsymbol{x}| \geq \gamma$, for some (unknown) gap parameter $\gamma > 0$, independent of \boldsymbol{x} .¹ As in the standard linear bandit setting (e.g., [5, 23, 12, 3, 13, 21, 30, 34, 16, 1], and references therein), the unknown vector \boldsymbol{u}_i determines the (average) behavior of user i. More concretely, upon receiving context vector \boldsymbol{x} , user i "reacts" by delivering a payoff value

$$a_i(\boldsymbol{x}) = \boldsymbol{u}_i^{\top} \boldsymbol{x} + \epsilon_i(\boldsymbol{x})$$

where $\epsilon_i(\boldsymbol{x})$ is a conditionally zero-mean and bounded variance noise term so that, conditioned on the past, the quantity $\boldsymbol{u}_i^{\top} \boldsymbol{x}$ is indeed the expected payoff observed at user *i* for context vector \boldsymbol{x} . Notice that the unknown parameter vector \boldsymbol{u}_i we associate with user *i* is supposed to be time invariant in this model.²

Since we are facing sequential decision settings where the learning system needs to continuously adapt to the newly received information provided by users, we assume that the learning process is broken up into a discrete sequence of rounds: In round $t = 1, 2, \ldots$, the learner receives a user index $i_t \in \mathcal{U}$ to serve content to, hence the user to serve may change at every round, though the same user can recur many times. We assume the sequence of users i_1, i_2, \ldots is determined by an exogenous process that places nonzero and independent probability to each user being the next one to serve. Together with i_t , the system receives in round t a set of feature vectors $C_{i_t} = \{ \boldsymbol{x}_{t,1}, \boldsymbol{x}_{t,2}, \dots, \boldsymbol{x}_{t,c_t} \} \subseteq \mathbb{R}^d$ encoding the content which is currently available for recommendation to user i_t . The learner is compelled to pick some $\bar{\boldsymbol{x}}_t = \boldsymbol{x}_{t,k_t} \in C_{i_t}$ to recommend to i_t , and then observes i_t 's feedback in the form of payoff $a_t \in \mathbb{R}$ whose (conditional) expectation is $\boldsymbol{u}_{i_t}^{\top} \bar{\boldsymbol{x}}_t$. The goal of the learning system is to maximize its total payoff $\sum_{t=1}^{T} a_t$ over T rounds. When the user feedback at our disposal is only the click/no-click behavior, the payoff a_t is naturally interpreted as a binary feedback, so that the quantity $\frac{\sum_{t=1}^{T} a_t}{T}$ becomes a clickthrough rate (CTR), where $a_t = 1$ if the recommended item was clicked by user i_t , and $a_t = 0$, otherwise. CTR is the measure of performance adopted by our comparative experiments in Section 5.

From a theoretical standpoint (Section 6), we are instead interested in bounding the cumulative *regret* achieved by our algorithms. More precisely, let the regret r_t of the learner at time t be the extent to which the average payoff of the best choice in hindsight at user i_t exceeds the average payoff of the algorithm's choice, i.e.,

$$r_t = \left(\max_{oldsymbol{x} \in C_{i_t}} oldsymbol{u}_{i_t}^{ op} oldsymbol{x}
ight) \!-\! oldsymbol{u}_{i_t}^{ op} oldsymbol{ar{x}}_t \;,$$

We are aimed at bounding with high probability the cumulative regret $\sum_{t=1}^{T} r_t$, the probability being over the noise variables $\epsilon_{i_t}(\bar{\boldsymbol{x}}_t)$, and any other possible source of randomness, including i_t – see Section 6.

¹As usual, this assumption may be relaxed by assuming the existence of two thresholds, one for the within-cluster distance of $u_i^{\top} x$ to $u_{i'}^{\top} x$, the other for the between-cluster distance.

 $^{^{2}}$ It would in fact be possible to lift this whole machinery to time-drifting user preferences by combining with known techniques (e.g., [11, 28]).

The kind of regret bound we would like to contrast to is one where the latent clustering structure over \mathcal{U} (w.r.t. the feature vectors \boldsymbol{x}) is somehow known beforehand (see Section 6 for details). When the content universe is large but known a priori, as is frequent in many collaborative filtering applications, it is often desirable to also group the items into clusters based on similarity of user preferences, i.e., two items are similar if they are preferred by many of the same users. This notion of "two-sided" clustering is well known in the literature; when the clustering process is simultaneously grouping users based on similarity at the item side and items based on similarity at the user side, it goes under the name of "co-clustering" (see, e.g., [14, 15]). Here, we consider a computationally more affordable notion of collaborate filtering based on adaptive two-sided clustering.

Unlike previous existing clustering techniques on bandits (e.g., [1, 29]), our clustering setting only applies to the case when the content universe is large but known a priori (yet, see the end of Section 4). Specifically, let the content universe be $\mathcal{I} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_{|\mathcal{I}|}\}$, and $P(\boldsymbol{x}_h) =$ $\{U_1(\boldsymbol{x}_h), U_2(\boldsymbol{x}_h), \ldots, U_{m(\boldsymbol{x}_h)}(\boldsymbol{x}_h)\}$ be the partition into clusters over the set of users \mathcal{U} induced by item \boldsymbol{x}_h . Then items $\boldsymbol{x}_h, \boldsymbol{x}_{h'} \in \mathcal{I}$ belong to the same cluster (over the set of items \mathcal{I}) if and only if they induce the same partition of the users, i.e., if $P(\boldsymbol{x}_h) = P(\boldsymbol{x}_{h'})$. We denote by g the number of distinct partitions so induced over \mathcal{U} by the items in \mathcal{I} , and work under the assumption that g is unknown but significantly smaller than $|\mathcal{I}|$. (Again, the assumption $g \ll |\mathcal{I}|$ is not strictly needed, but it both makes our algorithms effective and is expected to be satisfied in relevant practical scenarios.)

Finally, in all of the above, an important special case is when the items to be recommended do not possess specific features (or do not possess features having significant predictive power). In this case, it is common to resort to the more classical non-contextual stochastic multiarmed bandit setting (e.g., [6, 4]), which is recovered from the contextual framework by setting $d = |\mathcal{I}|$, and assuming the content universe ${\mathcal I}$ is made up of the d -dimensional vectors $e_h, h = 1, \ldots, d$, of the canonical basis of \mathbb{R}^d , As a consequence, the expected payoff of user i on item h is simply the *h*-th component of vector \boldsymbol{u}_i , and two users *i* and *i'* belong to the same cluster w.r.t. to h if the h-th component of u_i equals the *h*-th component of $u_{i'}$. Because the lack of useful annotation on data was an issue with all datasets at our disposal, it is this latter modeling assumption that motivates the algorithm we actually implemented for the experiments reported in Section 5.

3. RELATED WORK

Batch collaborative filtering neighborhood methods rely on finding similar groups of users and items to the target user-item pair, e.g., [33], and thus in effect rely on a dynamic form of grouping users and items. Collaborative Filteringbased methods have also been integrated with co-clustering techniques, whereby preferences in each co-cluster are modeled with simple statistics of the preference relations in the co-cluster, e.g., rating averages [18].

Beyond the general connection to co-clustering (e.g., [14, 15]), our paper is related to the research on multi-armed bandit algorithms for trading off exploration and exploitation through dynamic clustering. We are not aware of any specific piece of work that combines bandits with co-clustering based on the scheme of collaborative filtering; the papers which are most closely related to ours are [16, 27, 29, 8, 22, 1, 2]. In [16], the authors work under the assumption that users are defined using a feature vector, and try to learn a low-rank hidden subspace assuming that variation across users is low-rank. The paper combines low-rank matrix recovery with high-dimensional Gaussian Process Bandits, but it gives rise to algorithms which do not seem practical for sizeable problems. In [27], the authors analyze a non-contextual stochastic bandit problem where model parameters are assumed to be clustered in a few (unknown) types. Yet, the provided solutions are completely different from ours. The work [29] combines (k-means-like) online clustering with a contextual bandit setting, but clustering is only made at the user side. The paper [8] also relies on bandit clustering at the user side (as in [27, 29]), with an emphasis on diversifying recommendations to the same user over time. In [22], the authors propose cascading bandits of user behavior to identify the k most attractive items, and formulate it as a stochastic combinatorial partial monitoring problem. Finally, the algorithms in [1, 2, 26] can be seen as a special case of COFIBA when clustering is done only at the user side, under centralized [1, 26] or decentralized [2] environments.

Similar in spirit are also [7, 9, 10, 20]: In [7], the authors define a transfer learning problem within a stochastic multiarmed bandit setting, where a prior distribution is defined over the set of possible models over the tasks; in [9], the authors rely on clustering Markov Decision Processes based on their model parameter similarity. In [10], the authors discuss how to choose from n unknown distributions the kones whose means are largest by a certain metric; in [20] the authors study particle Thompson sampling with Rao-Blackwellization for online matrix factorization, exhibiting a regret bound in a very specific case of $n \times m$ rank-1 matrices. Yet, in none of above cases did the authors make a specific effort towards item-dependent clustering models applied to stochastic multi-armed bandits.

Further work includes [25, 32]. In [25], an ensemble of contextual bandits is used to address the cold-start problem in recommender systems. A similar approach is used in [32] to deal with cold-start in recommender systems but based on the probability matching paradigm in a parameter-free bandit strategy, which employs online bootstrap to derive the distribution of the estimated models. In contrast to our work, in neither [25] nor [32] are collaborative effects explicitly taken into account.

4. THE ALGORITHM

COFIBA, relies on upper-confidence-based tradeoffs between exploration and exploitation, combined with adaptive clustering procedures at both the user and the item sides. COFIBA stores in round t an estimate $\boldsymbol{w}_{i,t}$ of vector \boldsymbol{u}_i associated with user $i \in \mathcal{U}$. Vectors $\boldsymbol{w}_{i,t}$ are updated based on the payoff feedback, as in a standard linear least-squares approximation to the corresponding \boldsymbol{u}_i . Every user $i \in \mathcal{U}$ hosts such an algorithm which operates as a linear bandit algorithm (e.g., [12, 3, 1]) on the available content C_{i_t} . More specifically, $\boldsymbol{w}_{i,t-1}$ is determined by an inverse correlation matrix $M_{i,t-1}^{-1}$ subject to rank-one adjustments, and a vector $\boldsymbol{b}_{i,t-1}$ subject to additive updates. Matrices $M_{i,t}$ are initialized to the $d \times d$ identity matrix, and vectors $\boldsymbol{b}_{i,t}$ are initialized to the d-dimensional zero vector. Matrix $M_{i,t-1}^{-1}$ is also used to define an upper confidence bound $CB_{i,t-1}(\boldsymbol{x})$ in the approximation of $w_{i,t-1}$ to u_i along direction x. Based on the local information encoded in the weight vectors $\boldsymbol{w}_{i,t-1}$ and the confidence bounds $CB_{i,t-1}(\boldsymbol{x})$, the algorithm also maintains and updates a family of clusterings of the set of users \mathcal{U} , and a single clustering over the set of items \mathcal{I} . On both sides, such clusterings are represented through connected components of undirected graphs (this is in the same vein as in [1]), where nodes are either users or items. A pseudocode description of our algorithm is contained in Figures 1, 2, and 3, while Figure 4 illustrates the algorithm's behavior through a pictorial example.

At time t, COFIBA receives the index i_t of the current user to serve, along with the available item vectors $\boldsymbol{x}_{t,1},\ldots,\boldsymbol{x}_{t,c_t}$, and must select one among them. In order to do so, the algorithm computes the c_t neighborhood sets N_k , one per item $\boldsymbol{x}_{t,k} \in C_{i_t}$ based on the current aggregation of users (clusters "at the user side") w.r.t. item $x_{t,k}$. Set N_k should be regarded as the current approximation to the cluster (over the users) i_t belongs to when the clustering criterion is defined by item $\boldsymbol{x}_{t,k}$. Each neighborhood set then defines a compound weight vector $\bar{\boldsymbol{w}}_{N_k,t-1}$ (through the aggregation of the corresponding matrices $M_{i,t-1}$ and vectors $\boldsymbol{b}_{i,t-1}$ which, in turn, determines a compound confidence bound³ $CB_{N_k,t-1}(\boldsymbol{x}_{t,k})$. Vector $\bar{\boldsymbol{w}}_{N_k,t-1}$ and confidence bound $\operatorname{CB}_{N_k,t-1}(\boldsymbol{x}_{t,k})$ are combined through an upperconfidence exploration-exploitation scheme so as to commit to the specific item $\bar{\boldsymbol{x}}_t \in C_{i_t}$ for user i_t . Then, the payoff a_t is received, and the algorithm uses \bar{x}_t to update $M_{i_t,t-1}$ to $M_{i_t,t}$ and $\boldsymbol{b}_{i_t,t-1}$ to $\boldsymbol{b}_{i_t,t}$. Notice that the update is only performed at user i_t , though this will affect the calculation of neighborhood sets and compound vectors for other users in later rounds.

After receiving payoff a_t and computing $M_{i_t,t}$ and $b_{i_t,t}$, COFIBA updates the clusterings at the user side and the (unique) clustering at the item side. In round t, there are multiple graphs $G_{t,h}^{U} = (\mathcal{U}, E_{t,h}^{U})$ at the user side (hence many clusterings over \mathcal{U} , indexed by h), and a single graph $G_t^I = (\mathcal{I}, E_t^I)$ at the item side (hence a single clustering over \mathcal{I}). Each *clustering* at the user side corresponds to a single *cluster* at the item side, so that we have g_t *clusters* $I_{1,t}, \ldots, I_{g_t,t}$ over items and g_t clusterings over users – see Figure 4 for an example. On both user and item sides, updates take the form of edge deletions. Updates at the user side are only performed on the graph G_{t,\hat{h}_t}^U pointed to by the selected item $\bar{\boldsymbol{x}}_t = \boldsymbol{x}_{t,k_t}$. Updates at the item side are only made if it is likely that the neighborhoods of user i_t has significantly changed when considered w.r.t. two previously deemed similar items. Specifically, if item x_h was directly connected to item $\bar{\boldsymbol{x}}_t$ at the beginning of round t and, as a consequence of edge deletion at the user side, the set of users that are now likely to be close to i_t w.r.t. \boldsymbol{x}_h is no longer the same as the set of users that are likely to be close to i_t w.r.t. \bar{x}_t , then this is taken as a good indication that item \boldsymbol{x}_h is not inducing the same partition over users as $\bar{\boldsymbol{x}}_t$ does, hence edge $(\bar{\boldsymbol{x}}_t, \boldsymbol{x}_h)$ gets deleted. Notice that this need not imply that, as a result of this deletion, the two items are now belonging to different clusters over \mathcal{I} , since these two items may still be indirectly connected.

Input:

- Set of users $\mathcal{U} = \{1, \ldots, n\};$
- set of items $\mathcal{I} = \{ \boldsymbol{x}_1, \dots, \boldsymbol{x}_{|\mathcal{I}|} \} \subseteq \mathbb{R}^d;$
- exploration parameter $\alpha > 0$, and edge deletion parameter $\alpha_2 > 0$.

Init:

- $\boldsymbol{b}_{i,0} = \boldsymbol{0} \in \mathbb{R}^d$ and $M_{i,0} = I \in \mathbb{R}^{d \times d}, i = 1, \dots n;$
- User graph $G_{1,1}^U = (\mathcal{U}, E_{1,1}^U), G_{1,1}^U$ is connected over \mathcal{U} ;
- Number of user graphs $g_1 = 1$;
- No. of user clusters $m_{1,1}^U = 1$;
- Item clusters $\hat{I}_{1,1} = \mathcal{I}$, no. of item clusters $g_1 = 1$;
- Item graph $G_1^I = (\mathcal{I}, E_1^I), G_1^I$ is connected over \mathcal{I} .

for t = 1, 2, ..., T do Set

$$w_{i,t-1} = M_{i,t-1}^{-1} b_{i,t-1}, \qquad i = 1, \dots, n$$

Receive $i_t \in \mathcal{U}$, and get items $C_{i_t} = \{ \boldsymbol{x}_{t,1}, \dots, \boldsymbol{x}_{t,c_t} \} \subseteq \mathcal{I};$ For each $k = 1, ..., c_t$, determine which cluster (within the current user clustering w.r.t. $\boldsymbol{x}_{t,k}$) user i_t belongs to, and denote this cluster by N_k ; Compute, for $k = 1, \ldots, c_t$, aggregate quantities

$$\begin{split} \bar{M}_{N_k,t-1} &= I + \sum_{i \in N_k} (M_{i,t-1} - I), \\ \bar{\boldsymbol{b}}_{N_k,t-1} &= \sum_{i \in N_k} \boldsymbol{b}_{i,t-1}, \\ \bar{\boldsymbol{w}}_{N_k,t-1} &= \bar{M}_{N_k,t-1}^{-1} \bar{\boldsymbol{b}}_{N_k,t-1} ; \end{split}$$

Set

$$k_t = \operatorname*{argmax}_{k=1,...,c_t} \left(ar{oldsymbol{w}}_{N_k,t-1}^{ op} oldsymbol{x}_{t,k} + \operatorname{CB}_{N_k,t-1}(oldsymbol{x}_{t,k})
ight),$$

where $\operatorname{CB}_{N_k,t-1}(\boldsymbol{x}) = \alpha \sqrt{\boldsymbol{x}^\top \bar{M}_{N_k,t-1}^{-1} \boldsymbol{x} \log(t+1)};$ Set for brevity $\bar{\boldsymbol{x}}_t = \boldsymbol{x}_{t,k_t}$; Observe payoff $a_t \in \mathbb{R}$, and update weights $M_{i,t}$ and $b_{i,t}$ as follows:

- $M_{it,t} = M_{it,t-1} + \bar{\boldsymbol{x}}_t \bar{\boldsymbol{x}}_t^{\top},$
- $\boldsymbol{b}_{i_t,t} = \boldsymbol{b}_{i_t,t-1} + a_t \bar{\boldsymbol{x}}_t,$
- Set $M_{i,t} = M_{i,t-1}$, $\boldsymbol{b}_{i,t} = \boldsymbol{b}_{i,t-1}$ for all $i \neq i_t$,

Determine $\hat{h}_t \in \{1, \dots, g_t\}$ such that $k_t \in \hat{I}_{\hat{h}_t, t}$; Update user clusters at graph $G^U_{t, \hat{h}_t} = (\mathcal{U}, E^U_{t, \hat{h}_t})$ by performing the steps in Figure 2; For all $h \neq \hat{h}_t$, set $G_{t+1,h}^U = G_{t,h}^U$;

Update *item* clusters at graph $G_t^I = (\mathcal{I}, E_t^I)$ by performing the steps in Figure 3. end for

Figure 1: The COFIBA algorithm.

Update user clusters at graph G_{t,\hat{h}_t}^U as follows:

• Delete from E_{t,\hat{h}_t}^U all (i_t,j) such that $|\boldsymbol{w}_{i_t,t}^{\top} \bar{\boldsymbol{x}}_t - \boldsymbol{w}_{j,t}^{\top} \bar{\boldsymbol{x}}_t| > CB_{i_t,t}(\bar{\boldsymbol{x}}_t) + CB_{j,t}(\bar{\boldsymbol{x}}_t) ,$

where $CB_{i,t}(\boldsymbol{x}) = \alpha_2 \sqrt{\boldsymbol{x}^\top M_{i,t}^{-1} \boldsymbol{x} \log(t+1)};$

• Let E_{t+1,\hat{h}_t}^U be the resulting set of edges, set $G_{t+1,\widehat{h}_t}^U = (\mathcal{U}, E_{t+1,\widehat{h}_t}^U)$, and compute associated clus-
$$\begin{split} & \text{ters}\, \hat{U}_{1,t+1,\hat{h}_t}, \hat{U}_{2,t+1,\hat{h}_t}, \dots, \hat{U}_{m^U_{t+1,\hat{h}_t},t+1,\hat{h}_t} \text{ as the connected components of } G^U_{t+1,\hat{h}_t}. \end{split}$$

Figure 2: User cluster update in the COFIBA

 $^{^3\}mathrm{The}$ one given in Figure 1 is the confidence bound we use in our experiments. In fact, the theoretical counterpart to CB is significantly more involved, same efforts can also be found in order to close the gap, e.g., in [4, 1].

Update *item* clusters at graph G_t^I as follows:

• For all ℓ such that $(\bar{\boldsymbol{x}}_t, \boldsymbol{x}_\ell) \in E_t^I$ build neighborhood $N_{\ell,t+1}^U(i_t)$ as: $N_{\ell,t+1}^U(i_t) = \left\{ j : j \neq i_t, |\boldsymbol{w}_{i,t}^\top \boldsymbol{x}_\ell - \boldsymbol{w}_{i,t}^\top \boldsymbol{x}_\ell| \right\}$

$$\begin{aligned} {}_{+1}(i_t) &= \left\{ j \,:\, j \neq i_t \,, |\boldsymbol{w}_{i_t,t}^\top \boldsymbol{x}_\ell - \boldsymbol{w}_{j,t}^\top \boldsymbol{x}_\ell| \right. \\ &\leq {}_{\mathrm{CB}_{i_t,t}}(\boldsymbol{x}_\ell) + {}_{\mathrm{CB}_{j,t}}(\boldsymbol{x}_\ell) \right\}; \end{aligned}$$

- Delete from E_t^I all $(\bar{\boldsymbol{x}}_t, \boldsymbol{x}_\ell)$ such that $N_{\ell,t+1}^U(i_t) \neq N_{k_t,t+1}^U(i_t)$, where $N_{k_t,t+1}^U(i_t)$ is the *neighborhood* of node i_t w.r.t. graph G_{t+1,\widehat{h}_t}^U ;
- Let E_{t+1}^{I} be the resulting set of edges, set $G_{t+1}^{I} = (\mathcal{I}, E_{t+1}^{I})$, compute associated *item* clusters $\hat{I}_{1,t+1}, \hat{I}_{2,t+1}, \dots, \hat{I}_{g_{t+1},t+1}$ through the connected components of G_{t+1}^{I} ;
- For each new item cluster created, allocate a new connected graph over users representing a single (degenerate) cluster \mathcal{U} .

Figure 3: Item cluster update in the COFIBA

It is worth stressing that a naive implementation of COFIBA would require memory allocation for maintaining $|\mathcal{I}|$ -many *n*-node graphs, i.e., $\mathcal{O}(n^2 |\mathcal{I}|)$. Because this would be prohibitive even for moderately large sets of users, we make full usage of the approach of [1], where instead of starting off with complete graphs over users each time a new cluster over items is created, we randomly sparsify the complete graph by drawing an Erdos-Renyi initial graph, still retaining with high probability the underlying clusterings $\{U_1(\boldsymbol{x}_h), \ldots, U_{m(\boldsymbol{x}_h)}(\boldsymbol{x}_h)\}, h = 1, \ldots, |\mathcal{I}|, \text{ over users. This}$ works under the assumption that the latent clusters $U_i(\boldsymbol{x}_h)$ are not too small - see the argument in [1], where it is shown that in practice the initial graphs can have $\mathcal{O}(n \log n)$ edges instead of $\mathcal{O}(n^2)$. Moreover, because we modify the item graph by edge deletions only, one can show that with high probability (under the modeling assumptions of Section 2) the number q_t of clusters over items remains upper bounded by g throughout the run of COFIBA, so that the actual storage required by the algorithm is indeed $\mathcal{O}(ng\log n)$. This also brings a substantial saving in running time, since updating connected components scales with the number of edges of the involved graphs. It is this graph sparsification techniques that we used and tested along the way in our experimentation parts.

Finally, despite we have described in Section 2 a setting where \mathcal{I} and \mathcal{U} are known a priori (the analysis in Section 6 currently holds only in this scenario), nothing prevents in practice to adapt COFIBA to the case when new content or new users show up. This essentially amounts to adding new nodes to the graphs at either the item or the user side, by maintaining data-structures via dynamic memory allocation. In fact, this is precisely how we implemented our algorithm in the case of very big item or user sets (e.g., the Telefonica and the Avazu dataset in the next section).

5. EXPERIMENTS

We compared our algorithm to standard bandit baselines on three real-world datasets: one canonical benchmark dataset on news recommendations, one advertising dataset from a living production system, and one publicly available advertising dataset. In all cases, no features on the items have been used. We closely followed the same experimen-



Figure 4: In this example, $\mathcal{U} = \{1, \dots, 6\}$ and $\mathcal{I} =$ $\{x_1, \ldots, x_8\}$ (the items are depicted here as $1, 2, \ldots, 8$). (a) At the beginning we have $g_1 = 1$, with a single item cluster $\hat{I}_{1,1} = \mathcal{I}$ and, correspondingly, a single (degenerate) clustering over \mathcal{U} , made up of the unique cluster \mathcal{U} . (b) In round t we have the $g_t = 3$ item clusters $\hat{I}_{1,t} = \{ x_1, x_2 \}, \ \hat{I}_{2,t} = \{ x_3, x_4, x_5 \},$ $\hat{I}_{3,t} = \{ x_6, x_7, x_8 \}$. Corresponding to each one of them are the three clusterings over $\ensuremath{\mathcal{U}}$ depicted on the left, so that $m_{t,1}^U = 3$, $m_{t,2}^U = 2$, and $m_{t,3}^U = 4$. In this example, $i_t = 4$, and $\bar{x}_t = x_5$, hence $\hat{h}_t = 2$, and we focus on graph $G_{t,2}^U$, corresponding to user clustering $\{\{1, 2, 3\}, \{4, 5, 6\}\}$. Suppose in $G_{t,2}^U$ the only neighbors of user 4 are 5 and 6. When updating such user clustering, the algorithm considers therein edges (4,5)and (4, 6) to be candidates for elimination. Suppose edge (4,6) is eliminated, so that the new clustering over \mathcal{U} induced by the updated graph $G_{t+1,2}^U$ becomes $\{\{1, 2, 3\}, \{4, 5\}, \{6\}\}$. After user graph update, the algorithm considers the item graph update. Suppose x_5 is only connected to x_4 and x_3 in G_t^I , and that x_4 is not connected to x_3 , as depicted. Both edge (x_5, x_4) and edge (x_5, x_3) are candidates for elimination. The algorithm computes the neighborhood N of $i_t = 4$ according to $G_{t+1,2}^{U}$, and compares it to the the neighborhoods $N_{\ell,t+1}^U(i_t)$, for $\ell = 3, 4$. Assume $N \neq N_{3,t+1}^U(i_t)$, because the two neighborhoods of user 4 are now different, the algorithm deletes edge (x_5, x_3) from the item graph, splitting the item cluster $\{x_3, x_4, x_5\}$ into the two clusters $\{x_3\}$ and $\{x_4, x_5\}$, hence allocating a new cluster at the item side corresponding to a new degenerate clustering $\{\{1, 2, 3, 4, 5, 6\}\}$ at the user side. (c) The resulting clusterings at the beginning of round t + 1. (In this picture it is assumed that edge (x_5, x_4) was not deleted from the item graph at time t.)

tal setting as in previous work [12, 1], thereby evaluating prediction performance by click-through rate.

5.1 Datasets

Yahoo!. The first dataset we use for the evaluation is the

freely available benchmark dataset which was released in the "ICML 2012 Exploration & Exploitation Challenge"⁴. The aim of the challenge was to build state-of-the-art news article recommendation algorithms on Yahoo! data, by building an algorithm that learns efficiently a policy to serve news articles on a web site. The dataset is made up of random traffic records of user visits on the "Today Module" of Yahoo!, implying that both the visitors and the recommended news article are selected randomly. The available options (the items) correspond to a set of news articles available for recommendation, one being displayed in a small box on the visited web page. The aim is to recommend an interesting article to the user, whose interest in a given piece of news is asserted by a click on it. The data has 30 million visits over a two-week time stretch. Out of the logged information contained in each record, we used the user ID in the form of a 136-dimensional boolean vector containing his/her features (index i_t), the set of relevant news articles that the system can recommend from (set C_{i_t}); a randomly recommended article during the visit; a boolean value indicating whether the recommended article was clicked by the visiting user or not (payoff a_t). Because the displayed article is chosen uniformly at random from the candidate article pool, one can use an unbiased off-line evaluation method to compare bandit algorithms in a reliable way. We refer the reader to [1] for a more detailed description of how this dataset was collected and extracted. We picked the larger of the two datasets considered in [1], resulting in $n \approx 18K$ users, and d = 323 distinct items. The number of records ended up being 2.8M, out of which we took the first 300Kfor parameter tuning, and the rest for testing.

Telefonica. This dataset was obtained from Telefonica S.A., which is the number one Spanish broadband and telecommunications provider, with business units in Europe and South America. This data contains clicks on ads displayed to user on one of the websites that Telefonica operates on. The data were collected from the back-end server logs, and consist of two files: the first file contains the ads interactions (each record containing an impression timestamp, a user-ID, an action, the ad type, the order item ID, and the click timestamp); the second file contains the ads metadata as item-ID, type-ID, type, order-ID, creative type, mask, cost, creator-ID, transaction key, cap type. Overall, the number n of users was in the scale of millions, while the number d of items was approximately 300. The data contains 15M records, out of which we took the first 1, 5Mfor parameter tuning, and the rest for testing. Again, the only available payoffs are those associated with the items served by the system. Hence, in order to make the procedure be an effective estimator in a sequential decision process (e.g., [13, 17, 1, 23]), we simulated random choices by the system by generating the available item sets C_{i_t} as follows: At each round t, we stored the ad served to the current user i_t and the associated payoff value a_t (1 = "clicked", 0 = "not clicked"). Then we created C_{i_t} by including the served ad along with 9 extra items (hence $c_t = 10 \ \forall t$) which were drawn uniformly at random in such a way that, for any item $e_h \in \mathcal{I}$, if e_h occurs in some set C_{i_t} , this item will be the one served by the system 1/10 of the times. The random selection was done independent of the available payoff values a_t . All our experiments on this dataset were run on a machine with 64GB RAM and 32 Intel Xeon cores.

Avazu. This dataset was prepared by Avazu Inc⁵ which is a leading multinational corporation in the digital advertising business. The data was provided for the challenge to predict the click-through rate of impressions on mobile devices, i.e., whether a mobile ad will be clicked or not. The number of samples was around 40M, out of which we took the first 4M for parameter tuning, and the remaining for testing. Each line in the data file represents the event of an ad impression on the site or in a mobile application (app), along with additional context information. Again, payoff a_t is binary. The variables contained in the dataset for each sample are the following: ad-ID; timestamp (date and hour); click (boolean variable); device-ID; device IP; connection type; device type; ID of visited App/Website; category of visited App/Website; connection domain of visited App/Website; banner position; anonymized categorical fields (C1, C14-C21). We pre-processed the dataset as follows: we cleaned up the data by filtering out the records having missing feature values, and removed outliers. We identified the user with device-ID, if it is not null. The number of users on this dataset is in the scale of millions. Similar to the Telefonica dataset, we generated recommendation lists of length $c_t = 20$ for each distinct timestamp. We used the first 4M records for tuning parameters, and the remaining 36M for testing. All data were transferred to Amazon S3, and all jobs were run through the Amazon EC2 Web Service.

5.2 Algorithms

We compared COFIBA to a number of state-of-the-art bandit algorithms:

- LINUCB-ONE is a single instance of the UCB1 [6] algorithm, which is a very popular and established algorithm that has received a lot of attention in the research community over the past years;
- DYNUCB is the dynamic UCB algorithm of [29]. This algorithm adopts a "K-means"-like clustering technique so as to dynamically re-assign the clusters on the fly based on the changing contexts and user preferences over time;
- LINUCB-IND [1] is a set of independent UCB1 instances, one per user, which provides a fully personalized recommendation for each user;
- CLUB [1] is the state-of-the-art online clustering of bandits algorithm that dynamically cluster users based on the confidence ellipsoids of their models;
- LINUCB-V [4] is also a single instance of UCB1, but with a more sophisticated confidence bound; this algorithm turned out to be the winner of the "ICML 2012 Challenge" where the Yahoo! dataset originates from.

We tuned the optimal parameters in the training set with a standard grid search as indicated in [13, 1], and used the test set to evaluate the predictive performance of the algorithms. Since the system's recommendation need not coincide with the recommendation issued by the algorithms we tested, we only retained the records on which the two recommendations were indeed the same. Because records are discarded on

⁴https://explochallenge.inria.fr/category/challenge

⁵https://www.kaggle.com/c/avazu-ctr-prediction



Figure 5: Results on the Yahoo dataset.

the fly, the actual number T of retained records ("Rounds" in the plots of the next subsection) changes slightly across algorithms; T was around 70K for the Yahoo! data, 350K for the Telefonica data, and 900K for the Avazu data. All experimental results we report were averaged over 3 runs (but in fact the variance we observed across these runs was fairly small).

5.3 Results

Our results are summarized in Figures 5, 6, and 7. Further evidence is contained in Figure 8. In Figures 5–7, we plotted click-through rate ("CTR") vs. retained records so far ("Rounds"). All these experiments are aimed at testing the performance of the various bandit algorithms in terms of prediction performance, also in cold-start regimes (i.e., the first relatively small fraction of the time horizon in the xaxis). Our experimental setting is in line with previous ones (e.g., [12, 1]) and, by the way the data have been prepared, gives rise to a reliable estimation of actual CTR behavior under the same experimental conditions as in [12, 1]. Figure 8 is aimed at supporting the theoretical model of Section 2, by providing some evidence on the kind of clustering statistics produced by COFIBA at the end of its run.

Whereas the three datasets we took into consideration are all generated by real online web applications, it is worth pointing out that these datasets are indeed different in the way customers consume the associated content. Generally speaking, the longer the lifecycle of one item the fewer the items, the higher the chance that users with similar preferences will consume it, and hence the bigger the collaborative effects contained in the data. It is therefore reasonable to expect that our algorithm will be more effective in datasets where the collaborative effects are indeed strong.

The users in the Yahoo! data (Figure 5), are likely to span a wide range of demographic characteristics; on top of this, this dataset is derived from the consumption of news that are often interesting for large portions of these users and, as such, do not create strong polarization into subcommunities. This implies that more often than not, there are quite a few specific hot news that all users might express interest in, and it is natural to expect that these pieces of news are intended to reach a wide audience of consumers. Given this state of affairs, it is not surprising that on the Yahoo! dataset both LINUCB-ONE and LINUCB-V (serving the same news to all users) are already performing quite well, thereby making the clustering-of-users effort somewhat less useful. This also explains the poor performance of LINUCB-IND, which is



Figure 6: Results on the Telefonica dataset.



Figure 7: Results on the Avazu dataset.

not performing any clustering at all. Yet, even in this nontrivial case, COFIBA can still achieve a significant increased prediction accuracy compared, e.g., to CLUB, thereby suggesting that simultaneous clustering at both the user and the item (the news) sides might be an even more effective strategy to earn clicks in news recommendation systems.

Most of the users in the Telefonica data are from a diverse sample of people in Spain, and it is easy to imagine that this dataset spans a large number of communities across its population. Thus we can assume that collaborative effects will be much more evident, and that COFIBA will be able to leverage these effects efficiently. In this dataset, CLUB performs well in general, while DYNUCB deteriorates in the initial stage and catches-up later on. COFIBA seems to surpass all other algorithms, especially in the cold-start regime, all other algorithms being in the same ballpark as CLUB. Finally, the Avazu data is furnished from its professional digital advertising solution platform, where the customers click the ad impressions via the iOS/Android mobile apps or through websites, serving either the publisher or the advertiser which leads to a daily high volume internet traffic. In this dataset, neither LINUCB-ONE nor LINUCB-IND displayed a competitive cold-start performance. DYNUCB is underperforming throughout, while LINUCB-V demonstrates a relatively high CTR. CLUB is strong at the beginning, but then its CTR performance degrades. On the



Figure 8: A typical distribution of cluster sizes over users for the Yahoo dataset. Each bar plot corresponds to a cluster at the item side. We have 5 plots since this is the number of clusters over the items that COFIBA ended up with after sweeping once over this dataset in the run at hand. Each bar represents the fraction of users contained in the corresponding cluster. For instance, the first cluster over the items generated 16 clusters over the users (bar plot on top), with relative sizes 31%, 15%, 12%, etc. The second cluster over the items generated 10 clusters over the users (second bar plot from top) with relative sizes 61%, 12%, 9%, etc. The relative size of the 5 clusters over the items is as follows: 83%, 10%, 4%, 2%, and 1%, so that the clustering pattern depicted in the top plot applies to 83% of the items, the second one to 10% of the items, and so on.

other hand, COFIBA seems to work extremely well during the cold-start, and comparatively best in all later stages.

In Figure 8 we give a typical distribution of cluster sizes produced by COFIBA after at the end of its run.⁶ The emerging pattern is always the same: we have few clusters over the items with very unbalanced sizes and, corresponding to each item cluster, we have few clusters over the users, again with very unbalanced sizes. This recurring pattern is in fact the motivation behind our theoretical assumptions (Section 2), and a property of data that the COFIBA algorithm can provably take advantage of (Section 6). These bar plots, combined with the comparatively good performance of COFIBA, suggest that our datasets do actually possess clusterability properties at both sides.

To summarize, despite the differences in the three

datasets, the experimental evidence we collected on them is quite consistent, in that in all the three cases COFIBA significantly outperforms all other competing methods we tested. This is especially noticeable during the cold-start period, but the same relative behavior essentially shows up during the whole time window of our experiments. COFIBA is a bit involved to implement, as contrasted to its competitors, and is also somewhat slower to run (unsurprisingly slower than, say, LINUCB-ONE and LINUCB-IND). On the other hand, COFIBA is far more effective in exploiting the collaborative effects embedded in the data, and still amenable to be run on large datasets.

6. REGRET ANALYSIS

The following theorem is the theoretical guarantee of COFIBA, where we relate the cumulative regret of COFIBA to the clustering structure of users \mathcal{U} w.r.t. items \mathcal{I} . For simplicity of presentation, we formulate our result in the one-hot encoding case, where $\boldsymbol{u}_i \in \mathbb{R}^d$, $i = 1, \ldots, n$, and $\mathcal{I} = \{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d\}$. In fact, a more general statement can be proven which holds in the case when \mathcal{I} is a generic set of feature vectors $\mathcal{I} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{|\mathcal{I}|}\}$, and the regret bound depends on the geometric properties of such vectors.⁷

In order to obtain a provable advantage from our clusterability assumptions, extra conditions are needed on the way i_t and C_{i_t} are generated. The clusterability assumptions we can naturally take advantage of are those where, for most partitions $P(e_h)$, the relative sizes of clusters over users are highly unbalanced. Translated into more practical terms, cluster unbalancedness amounts to saying that the universe of items \mathcal{I} tends to influence users so as to determine a small number of major common behaviors (which need neither be the same nor involve the same users across items), along with a number of minor ones. As we saw in our experiments, this seems like a frequent behavior of users in some practical scenarios.

THEOREM 1. Let the COFIBA algorithm of Figure 1 be run on a set of users $\mathcal{U} = \{1, \ldots, n\}$ with associated profile vectors $\boldsymbol{u}_1,\ldots,\boldsymbol{u}_n\in\mathbb{R}^d$, and set of items $\mathcal{I}=\{\boldsymbol{e}_1,\ldots,\boldsymbol{e}_d\}$ such that the h-th induced partition $P(e_h)$ over \mathcal{U} is made up of m_h clusters of cardinality $v_{h,1}, v_{h,2}, \ldots, v_{h,m_h}$, respectively. Moreover, let g be the number of distinct partitions so obtained. At each round t, let i_t be generated uniformly at random⁸ from \mathcal{U} . Once i_t is selected, the number c_t of items in C_{i_t} is generated arbitrarily as a function of past indices i_1, \ldots, i_{t-1} , payoffs a_1, \ldots, a_{t-1} , and sets $C_{i_1}, \ldots, C_{i_{t-1}}$, as well as the current index i_t . Then the sequence of items in C_{i_t} is generated i.i.d. (conditioned on i_t , c_t and all past indices i_1, \ldots, i_{t-1} , payoffs a_1, \ldots, a_{t-1} , and sets $C_{i_1}, \ldots, C_{i_{t-1}}$) according to a given but unknown distribution \mathcal{D} over \mathcal{I} . Let payoff a_t lie in the interval [-1, 1], and be generated as described in Section 2 so that, conditioned on history, the expectation of a_t is $\boldsymbol{u}_{i_t}^{\top} \bar{\boldsymbol{x}}_t$. Finally, let parameters α and α_2 be suitable functions of $\log(1/\delta)$. If $c_t \leq c \ \forall t \ then, \ as \ T \ grows \ large, \ with \ probability \ at \ least$

 $^{^6}$ Without loss of generality, we take the first Yahoo dataset to provide statistics, for similar shapes of the bar plots can be established for the remaining ones.

 $^{^{7}\}mathrm{In}$ addition, the function CB should be modified so as to incorporate these properties.

⁸Any distribution having positive probability on each $i \in \mathcal{U}$ would suffice here.

 $1 - \delta$ the cumulative regret satisfies⁹

$$\sum_{t=1}^{T} r_t = \widetilde{\mathcal{O}}\left(\left(\mathbb{E}[S] + \sqrt{c\sqrt{mn}\operatorname{VAR}(S)} + 1\right)\sqrt{\frac{dT}{n}}\right) ,$$

where $S = S(h) = \sum_{j=1}^{m_h} \sqrt{v_{h,j}}$, h is a random index such that $\mathbf{e}_h \sim \mathcal{D}$, and $\mathbb{E}[\cdot]$ and $\operatorname{VAR}(\cdot)$ denote, respectively, the expectation and the variance w.r.t. this random index.

To get a feeling of how big (or small) $\mathbb{E}[S]$ and $\operatorname{VAR}[S]$ can be, let us consider the case where each partition over users has a single big cluster and a number of small ones. To make it clear, consider the extreme scenario where each $P(e_h)$ has one cluster of size $v_{h,1} = n - (m - 1)$, and m - 1 clusters of size $v_{h,j} = 1$, with $m < \sqrt{n}$. Then it is easy to see that $\mathbb{E}[S] = \sqrt{n - (m - 1)} + m - 1$, and VAR(S) = 0, so that the resulting regret bound essentially becomes $\mathcal{O}(\sqrt{dT})$, which is the standard regret bound one achieves for learning a single d-dimensional user (aka, the standard noncontextual bandit bound with d actions and no gap assumptions among them). At the other extreme lies the case when each partition $P(e_h)$ has n-many clusters, so that $\mathbb{E}[S] = n$, $\operatorname{VAR}(S) = 0$, and the resulting bound is $\widetilde{\mathcal{O}}(\sqrt{dnT})$. Looser upper bounds can be achieved in the case when VAR(S) > 0, where also the interplay with c starts becoming relevant. Finally, observe that the number g of distinct partitions influences the bound only indirectly through VAR(S). Yet, it is worth repeating here that q plays a crucial role in the computational (both time and space) complexity of the whole procedure.

PROOF OF THEOREM 1. The proof sketch builds on the analysis in [1]. Let the true underlying clusters over the users be $V_{h,1}, V_{h,2}, ..., V_{h,m_h}$, with $|V_{h,j}| = v_{h,j}$. In [1], the authors show that, because each user i has probability 1/nto be the one served in round t, we have, with high probability, $\boldsymbol{w}_{i,t} \rightarrow \boldsymbol{u}_i$ for all *i*, as *t* grows large. Moreover, because of the gap assumption involving parameter γ , all edges connecting users belonging to different clusters at the user side will eventually be deleted (again, with high probability), after each user i is served at least $O(\frac{1}{\gamma^2})$ times. By the way edges are disconnected at the item side, the above is essentially independent (up to log factors due to union bounds) of which graph at the user side we are referring to. In turn, this entails that the current user clusters encoded by the connected components of graph $G_{t,h}^U$ will eventually converge to the m_h true user clusters (again, independent of h, up to log factors), so that the aggregate weight vectors $ar{m{w}}_{N_k,t-1}$ computed by the algorithm for trading off exploration vs. exploitation in round t will essentially converge to \boldsymbol{u}_{i_t} at a rate of the form¹⁰

$$\mathbb{E}\left[\frac{1}{\sqrt{1+T_{h_t,j_t,t-1}/d}}\right],\qquad(1)$$

where h_t is the index of the true cluster over *items* that \bar{x}_t belongs to, j_t is the index of the true cluster over *users* that i_t belongs to (according to the partition of \mathcal{U} determined

by h_t), $T_{h_t, j_t, t-1}$ is the number of rounds so far where we happened to "hit" cluster V_{h_t, j_t} , i.e.,

$$T_{h_t, j_t, t-1} = |\{s \le t - 1 : i_s \in V_{h_t, j_t}\}|,$$

and the expectation is w.r.t. both the (uniform) distribution of i_t , and distribution \mathcal{D} generating the items in C_{i_t} , conditioned on all past events. Since, by the Azuma-Hoeffding inequality, $T_{h_t,j_t,t-1}$ concentrates as

$$T_{h_t,j_t,t-1} \approx \frac{t-1}{n} v_{h_t,j_t} ,$$

we have

(1)
$$\approx \mathbb{E}_{\mathcal{D}}\left[\sum_{j=1}^{m_{h_t}} \frac{v_{h_t,j}}{n} \frac{1}{\sqrt{1 + \frac{t-1}{d\,n}\,v_{h_t,j}}}\right]$$

It is the latter expression that rules the cumulative regret of COFIBA in that, up to log factors:

$$\sum_{t=1}^{T} r_t \approx \sum_{t=1}^{T} \mathbb{E}_{\mathcal{D}} \left[\sum_{j=1}^{m_{h_t}} \frac{v_{h_t,j}}{n} \frac{1}{\sqrt{1 + \frac{t-1}{dn} v_{h_t,j}}} \right] .$$
 (2)

Eq. (2) is essentially (up to log factors and omitted additive terms) the regret bound one would obtain by knowning beforehand the latent clustering structure over \mathcal{U} .

Because $h_t \in C_{i_t}$ is itself a function of the items in C_{i_t} , we can eliminate the dependence on h_t by the following simple stratification argument. First of all, notice that

$$\sum_{j=1}^{n_{h_t}} \frac{v_{h_t,j}}{n} \frac{1}{\sqrt{1 + \frac{t-1}{d n} v_{h_t,j}}} \approx \sqrt{\frac{d}{nt}} \sum_{j=1}^{m_{h_t}} \sqrt{v_{h_t,j}} \; .$$

Then, we set for brevity $S(h) = \sum_{j=1}^{m_h} \sqrt{v_{h,j}}$, and let $h_{t,k}$ be the index of the true cluster over items that $\boldsymbol{x}_{t,k}$ belongs to (recall that $h_{t,k}$ is a random variable since so is $\boldsymbol{x}_{t,k}$). Since $S(h_{t,k}) \leq \sqrt{mn}$, a standard argument shows that

$$\mathbb{E}_{\mathcal{D}}[S(h_t)] \leq \mathbb{E}_{\mathcal{D}}\left[\max_{k=1,\dots,c_t} S(h_{t,k})\right]$$
$$\leq \mathbb{E}_{\mathcal{D}}[S(h_{t,1})] + \sqrt{c\sqrt{mn}\operatorname{VAR}_{\mathcal{D}}(S(h_{t,1}))} + 1 ,$$

so that, after some overapproximations, we conclude that $\sum_{t=1}^{T} r_t$ is upper bounded with high probability by

$$\widetilde{\mathcal{O}}\left(\left(\mathbb{E}_{\mathcal{D}}[S(h)] + \sqrt{c\sqrt{mn}\operatorname{VAR}_{\mathcal{D}}(S(h))} + 1\right)\sqrt{\frac{dT}{n}}\right) ,$$

the expectation and the variance being over the random index h such that $e_h \sim \mathcal{D}$. \Box

7. CONCLUSIONS

We have initiated an investigation of collaborative filtering bandit algorithms operating in relevant scenarios where multiple users can be grouped by behavior similarity in different ways w.r.t. items and, in turn, the universe of items can possibly be grouped by the similarity of clusterings they induce over users. We carried out an extensive experimental comparison with very encouraging results, and have also given a regret analysis which operates in a simplified scenario. Our algorithm can in principle be modified so as to be combined with any standard clustering (or co-clustering) technique. However, one advantage of encoding clusters as connected components of graphs (at least at the user side) is

⁹The $\widetilde{\mathcal{O}}$ -notation hides logarithmic factors in $n, m, g, T, d, 1/\delta$, as well as terms which are independent of T.

¹⁰Because $\mathcal{I} = \{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d\}$, the minimal eigenvalue λ of the process correlation matrix $\mathbb{E}[X X^{\top}]$ in [1] is here 1/d. Moreover, compared to [1], we do not strive to capture the geometry of the user vectors \boldsymbol{u}_i in the regret bound, hence we do not have the extra \sqrt{m} factor occurring in their bound.

that we are quite effective in tackling the so-called *cold start* problem, for the newly served users are more likely to be connected to the old ones, which makes COFIBA in a position to automatically propagate information from the old users to the new ones through the aggregate vectors $\bar{\boldsymbol{w}}_{N_k,t}$. In fact, so far we have not seen any other way of adaptively clustering users and items which is computationally affordable on sizeable datasets and, at the same time, amenable to a regret analysis that takes advantage of the clustering assumption.

All our experiments have been conducted in the setup of one-hot encoding, since the datasets at our disposal did not come with reliable/useful annotations on data. Yet, the algorithm we presented can clearly work when the items are accompanied by (numerical) features. One direction of our future research is to compensate for the lack of features in the data by first *inferring* features during an initial training phase through standard matrix factorization techniques, and subsequently applying our algorithm to a universe of items \mathcal{I} described through such inferred features. Another line of experimental research would be to combine different bandit algorithms (possibly at different stages of the learning process) so as to roughly get the best of all of them in all stages. This would be somewhat similar to the metabandit construction described in [25]. Another one would be to combine with matrix factorization techniques as in, e.g., [20].

8. ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for their helpful and constructive comments. The first author thanks the support from MIUR and QCRI-HBKU. Also, the first and the third author acknowledge the support from Amazon AWS Award in Machine Learning Research Grant. The work leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/ 2007-2013) under CrowdRec Grant Agreement n° 610594.

9. **REFERENCES**

- [1] C. Gentile, S. Li, and G. Zappella. Online clustering of bandits. In *ICML*, 2014.
- [2] N. Korda, B. Szorenyi, and S. Li, Distributed Clustering of Linear Bandits in Peer to Peer Networks. In *ICML*, 2016.
- [3] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. 2011.
- [4] J.-Y. Audibert, R. Munos, and C. Szepesvári. Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.
- P. Auer. Using confidence bounds for exploration-exploitation trade-offs. Journal of Machine Learning Research, 3:397–422, 2002.
- [6] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 2001.
- [7] M. G. Azar, A. Lazaric, and E. Brunskill. Sequential transfer in multi-armed bandit with finite set of models. In *NIPS*, pages 2220–2228, 2013.
- [8] G. Bresler, G. Chen, and S. D. A latent source model for online collaborative filtering. In *NIPS*. MIT Press, 2014.
- [9] E. Brunskill and L. Li. Sample complexity of multi-task reinforcement learning. In UAI, 2013.
- [10] W. Cao, J. Li, Y. Tao, and Z. Li. On top-k selection in multi-armed bandits and hidden bipartite graphs. In NIPS, 2015.

- [11] G. Cavallanti, N. Cesa-Bianchi, and C. Gentile. Tracking the best hyperplane with a simple budget perceptron. *Machine Learning*, 69/2:143–167, 2007.
- [12] W. Chu, L. Li, L. Reyzin, and R. E. Schapire. Contextual bandits with linear payoff functions. In AISTATS, 2011.
- [13] K. Crammer and C. Gentile. Multiclass classification with bandit feedback using adaptive regularization. In *ICML*, 2011.
- [14] I. S. Dhillon. Co-clustering documents and words using bipartite spectral graph partitioning. In 7th KDD, pages 269–274. ACM, 2001.
- [15] I. S. Dhillon, S. Mallela, and D. S. Modha. Information-theoretic co-clustering. In 9th KDD, pages 89–98, New York, NY, USA, 2003. ACM.
- [16] J. Djolonga, A. Krause, and V. Cevher. High-dimensional gaussian process bandits. In NIPS, pages 1025–1033, 2013.
- [17] M. Dudik, D. Erhan, J. Langford, and L. Li. Sample-efficient nonstationary-policy evaluation for contextual bandits. In UAI, 2012.
- [18] T. George and S. Merugu. A scalable collaborative filtering framework based on co-clustering. In 5th ICDM, pages 625–628. IEEE Computer Society, 2005.
- [19] F. Hao, S. Li, G. Min, H. Kim, S. Yau, and L. Yang, An Efficient Approach to Generating Location-Sensitive Recommendations in Ad-hoc Social Network Environments. *IEEE Transactions on Services Computing*, 8:3, pp. 520–533, 2015.
- [20] J. Kawale, H. Bui, B. Kveton, L. Thanh, and S. Chawla. Efficient thompson sampling for online matrix-factorization recommendation. In NIPS, 2015.
- [21] A. Krause and C. Ong. Contextual gaussian process bandit optimization. In 25th NIPS, 2011.
- [22] B. Kveton, C. Szepesvari, Z. Wen, and A. Ashkan. Cascading bandits: Learning to rank in the cascade model. In *ICML*, 2015.
- [23] L. Li, W. Chu, J. Langford, and R. Schapire. A contextual-bandit approach to personalized news article recommendation. In WWW, pages 661–670, 2010.
- [24] S. Li, F. Hao, M. Li, and H.C. Kim. Medicine Rating Prediction and Recommendation in Mobile Social Networks. In International Conference on Green, Pervasive and Cloud Computing, 2013.
- [25] L. Tang, Y. Jiang, L. Li, and T. Li. Ensemble contextual bandits for personalized recommendation. In *RecSys*, 2014.
- [26] S. Li, C. Gentile, and A. Karatzoglou. Graph clustering bandits for recommendation. arXiv:1605.00596.
- [27] O. Maillard and S. Mannor. Latent bandits. In ICML, 2014.
- [28] E. Moroshko, N. Vaits, and K. Crammer. Second-order non-stationary online learning for regression. *Journal of Machine Learning Research*, 16:1481–1517, 2015.
- [29] T. T. Nguyen and H. W. Lauw. Dynamic clustering of contextual multi-armed bandits. In 23rd CIKM, pages 1959–1962. ACM, 2014.
- [30] Y. Seldin, P. Auer, F. Laviolette, J. Shawe-Taylor, and R. Ortner. Pac-bayesian analysis of contextual bandits. In *NIPS*, pages 1683–1691, 2011.
- [31] I. Sutskever, R. Salakhutdinov, and J. Tenenbaum. Modelling relational data using bayesian clustered tensor factorization. In *NIPS*, pages 1821–1828. MIT Press, 2009.
- [32] L. Tang, Y. Jiang, L. Li, C. Zeng, and T. Li. Personalized recommendation via parameter-free contextual bandits. In *SIGIR*. ACM, 2015.
- [33] K. Verstrepen and B. Goethals. Unifying nearest neighbors collaborative filtering. In *RecSys*, 2014.
- [34] Y. Yue, S. A. Hong, and C. Guestrin. Hierarchical exploration for accelerating contextual bandits. In *ICML*, 2012.