

Social Collaborative Retrieval

Ko-Jen Hsiao
University of Michigan
Ann Arbor, MI, USA 48109
coolmark@umich.edu

Alex Kulesza
University of Michigan
Ann Arbor, MI, USA 48109
kulesza@umich.edu

Alfred O. Hero III
University of Michigan
Ann Arbor, MI, USA 48109
hero@umich.edu

ABSTRACT

Socially-based recommendation systems have recently attracted significant interest, and a number of studies have shown that social information can dramatically improve a system’s predictions of user interests. Meanwhile, there are now many potential applications that involve aspects of both recommendation and information retrieval, and the task of *collaborative retrieval*—a combination of these two traditional problems—has recently been introduced. Successful collaborative retrieval requires overcoming severe data sparsity, making additional sources of information, such as social graphs, particularly valuable. In this paper we propose a new model for collaborative retrieval, and show that our algorithm outperforms current state-of-the-art approaches by incorporating information from social networks. We also provide empirical analyses of the ways in which cultural interests propagate along a social graph using a real-world music dataset.

Categories and Subject Descriptors

H.3.3 [Information storage and retrieval]: Information Search and Retrieval; I.2.6 [Artificial Intelligence]: Learning

General Terms

Algorithms, Experimentation

Keywords

collaborative filtering, information retrieval, social networks

1. INTRODUCTION

Collaborative filtering (CF) and related recommendation techniques, which aim to automatically predict the interests of users and make personal recommendations of music, movies, products, or other items, have been both intensively studied by researchers and successfully deployed in industry

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 the author(s) 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.
WSDM '14, February 24–28, 2014, New York, New York, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2351-2/14/02 ...\$15.00.

<http://>—enter the whole DOI string from rightsreview form confirmation.

during the past decade [3, 9, 18, 21]. Recently, Weston et al. [24] proposed extending CF to a setting termed collaborative retrieval (CR), in which recommendations are made with respect to a particular *query context*; for instance, a user might be looking for music in a particular genre or movies similar to a recent favorite. In these situations, otherwise accurate recommendations will become irrelevant. Similarly, a shopping website might want to deliver a list of recommended items to a user based on their browsing history. In this case the recently viewed pages act as a sort of query, and we would like recommendations that are specific both to the query and to the user. Weston et al. [24] proposed the first algorithm to solve CR problems, called latent collaborative retrieval (LCR).

However, several important issues remain. While it is well-known that CF models typically contend with data sparsity problems since a large number of user-item pairs must be judged for compatibility based on a relatively small dataset, CR problems suffer even more severely from sparsity, since the range of possible queries multiplies the number of judgments to be made. Viewed as a matrix completion problem, traditional CF requires filling out a user \times item matrix, where each entry indicates the relevance of a specific item to a specific user; this is often done using techniques like singular value decomposition and non-negative matrix factorization (NMF) [19]. In the same light, CR can be seen as a tensor completion problem, where the goal is to fill out a much larger query \times user \times item tensor. Recent work has begun to extend matrix factorization techniques for CF to tensor models of this type [6, 22, 25, 26]; however, these methods typically do not accommodate the ranking losses used for CR, and sparsity remains an important concern. In this work, we propose to deal with data sparsity in CR by incorporating an additional (but often readily available) source of information: a social graph.

In recent years social networking sites have become extremely popular, producing significant insights into the ways in which people connect themselves and interact. Information derived from these networks can be used to help address the sparsity problem faced by recommender systems, for instance by propagating information from a user’s social connections to fill in gaps in the recommendation matrix. A variety of CF models augmented with social information have recently been proposed [5, 7, 8, 10, 12, 16]; these include state-of-the-art methods like Collaborative Topic Regression with social matrix factorization [16], which is based on LDA [2], and Probabilistic Matrix Factorization (PMF) [10, 17]. There has also been interest in so-called *trust-aware*

recommendation methods [1, 11, 13, 14, 15], which are similar in spirit but inherently limited compared with using real social networks [12]. However, social information has not yet been made to work for collaborative retrieval, which arguably stands to benefit even more. In this paper we set out to fill this gap.

We propose an approach we call social collaborative retrieval (SCR), which builds on the latent collaborative retrieval (LCR) model of Weston et al. [24] by integrating social networking data. Notice that this work is different from researches in collaborative information seeking (CIS). CIS is a popular field of research that studies methods and systems for people working in collaborative group for information seeking. In [20], CIS is defined as "a process of collaboratively seeking information that is defined explicitly among the participants, interactive, and mutually beneficial." CIS can be viewed as an active process and is quite different from collaborative filtering. While our work is more related to collaborative filtering, our SCR model connects users via their social network in addition to their passive involvements (e.g., listening to similar music artists).

As in LCR, our algorithm sets out to optimize the top-ranked items retrieved for a given user and query, but we incorporate a regularization penalty that encourages the low dimensional embedding of each user to be similar to those of the user's nearest social neighbors. On a collaborative retrieval task using real-world artist ratings from the Last.fm music dataset, our proposed algorithm significantly outperforms LCR, as well as baseline CF methods based on non-negative matrix factorization and singular value decomposition, particularly when a smaller training set leads to increased data sparsity.

The rest of this paper is organized as follows. In Section 2 we provide a brief introduction to latent collaborative retrieval (LCR) [24], followed by the details of our proposed SCR algorithm. Section 3 presents an empirical analysis of the Last.fm social networking dataset, and we present experimental results evaluating the performance of SCR in Section 4.

2. COLLABORATIVE RETRIEVAL

The goal of collaborative retrieval is to produce a ranked list of items that are of interest to a particular user given a particular query. While one approach might be to simply filter a set of unconstrained CF recommendations for the specified user using the query—or, conversely, to filter a set of generic search results for the query using the user's profile—these pipeline approaches fail to account for interactions between the query and the user. For instance, two users might have very different interpretations of the query "jazz", despite having broadly similar preferences among artists. CR attempts to obtain more informative results through a unified approach.

We thus formalize the problem by defining a single score function $f(q, u, a)$ to represent the relevance of a given item a with respect to both a query q and a user u . If we enumerate all users, queries, and items, we can think of this score function as specifying the values of a rating tensor $\mathbf{R} \in \mathbb{R}^{|\mathcal{Q}| \times |\mathcal{U}| \times |\mathcal{A}|}$, where \mathcal{Q} is the set of all queries, \mathcal{U} is the set of users, and \mathcal{A} is the set of items. However, in practice we usually only care about the top k items retrieved (for some small constant k) for a given user and query, and our evaluation metrics will share this property as well. (We dis-

cuss specific error measures in Section 2.3.) Thus, learning a score function that can correctly *rank* items for a given user-query pair is more important than learning one which can correctly approximate the full set of entries in \mathbf{R} . The objectives that we use to learn the parameters of the score function will therefore involve a measure of error on such top- k ranked lists.

We next briefly review the existing latent collaborative retrieval model for this problem, and then introduce our model using social information. Finally, we discuss the optimization needed to learn the parameters of the model.

2.1 Latent collaborative retrieval

Latent collaborative retrieval (LCR) [24] was the first algorithm proposed to solve collaborative retrieval problems. The central idea is to embed users, queries, and items in a common n -dimensional space in which they can be compared using linear operations. (n is a hyperparameter that is typically tuned on a validation set.) Formally, LCR is parameterized by matrices $S \in \mathbb{R}^{n \times |\mathcal{Q}|}$, $V \in \mathbb{R}^{n \times |\mathcal{U}|}$, and $T \in \mathbb{R}^{n \times |\mathcal{A}|}$, which give the low-dimensional representations of queries, users, and items, respectively. Additionally, for each user u a matrix $U_u \in \mathbb{R}^{n \times n}$ encodes the user-specific relationship between queries and items. The scoring function f is then defined as

$$f(q, u, a) = S_q^\top U_u T_a + V_u^\top T_a, \quad (1)$$

where S_q is the column of S corresponding to query q , T_a is the column of T corresponding to item a , and V_u is the column of V corresponding to user u . Intuitively, the first term in Equation 1 measures the similarity between the query and the item in the low-dimensional space under a linear transformation that is dependent on the user. The second term is independent of the query and can be viewed as a bias term which models user preferences for different items. Since for a given instance of a CR task the query and user are held fixed, there is no need for the scoring function to include a term like $S_q^\top V_u$, which would measure the compatibility of a user and a query. However, interactions between the user and the query that pertain to actual item recommendations can be expressed in the first term. If there are significant non-user-specific aspects of the compatibility between queries and items (i.e., a $S_q^\top T_a$ term), these can simply be absorbed into the first term and need not appear separately.

The parameters of the LCR scoring function are learned by optimizing a chosen error metric over a training set; we discuss some such metrics and other details in Section 2.3.

To aid in generalization, and to avoid the potentially prohibitive enumeration of queries, Equation 1 can be generalized using features. In this case $\Phi_Q(q)$, $\Phi_U(u)$ and $\Phi_A(a)$ are vector-valued feature maps for queries, users, and items, respectively, and S , T , and V are linear maps from feature vectors to the embedded n -dimensional space. The feature-based scoring function is given by

$$f(q, u, a) = \Phi_Q(q)^\top S^\top U_u T \Phi_A(a) + \Phi_U(u)^\top V^\top T \Phi_A(a). \quad (2)$$

If the feature maps are simple characteristic vectors, with a unique feature for each query, user, and item, then we recover the simpler form of Equation 1. Features of this type can also be used for content-based recommendation models (see [24]). For our purposes, we simply note that this feature-based formulation can be easily extended to SCR,

but for simplicity we focus on the form in Equation 1 going forward.

2.2 Social collaborative retrieval

In the real world, people often turn to their friends for recommendations of musics, movies, or products. Here, we apply this intuition to improve the performance of CR techniques on tasks where social information is available. Our approach, which we refer to as social collaborative retrieval (SCR), learns a scoring function using a combination of *behavioral* and *relational* error measures.

Behavioral measures encourage the model to respect implicit similarities between users, items, and queries that are revealed by the training data. For instance, the preferences of one user may be useful for recommending items to another user if the two users have expressed similar preferences in the past. This is the traditional mode of operation for collaborative filtering, as well as for CR.

Relational measures, on the other hand, take account of explicitly labeled connections that (hopefully) reveal underlying similarities. In this work, we employ a relational measure that encourages the scoring function to be smooth with respect to a social graph; that is, we assume that users who are social neighbors should, on average, have more similar preferences than those who are not. (We validate this assumption empirically in Section 3.) The hope is that this relational measure term provides complementary guidance to the system when little is known about the behavior of a user.

For simplicity, and to make the fairest comparisons later, we use the same parameterization of the scoring function as LCR (Equation 1); we have n -dimensional representations of queries, users, and items in matrices S , V , and T , respectively, as well as user-specific transformations $U_u \in \mathbb{R}^{n \times n}$. We additionally assume that a social graph G is available, where $G(i, j) = 1$ whenever users i and j are linked, and $G(i, j) = 0$ otherwise. We will sometimes refer to users who are linked as “friends”, though of course the social graph may encode relations with varying semantics. To bias the preferences of friends toward each other, SCR introduces the social error measure

$$\text{err}_{\text{social}}(V, G) = \sum_{i,j,G(i,j)=1} \|\sigma(V_i^T V_j) - 1\|^2, \quad (3)$$

where $\sigma(\cdot)$ is the sigmoid function

$$\sigma(x) = \frac{1}{1 + e^{-cx}} \quad (4)$$

and c is a hyperparameter. This measure can be seen as a regularization penalty that is minimized when friends have identical, high-norm representations in V . Notice that we do not penalize similarity among non-friends, since users may have similar tastes even though they are not friends. Importantly, although we encourage friends to have similar representations V_u , we do not introduce such regularization for U_u matrices, as this would tend to force friends to always receive the same results. Intuitively, we expect that friends are likely to have similar taste in items, but we allow each their own particular querying “style”.

Combining the relational measure in Equation 3 with a behavioral measure $\text{err}_{\text{behavior}}$ that depends on the scoring function f and the training set X yields the SCR learning

objective to be minimized:

$$\text{err}_{\text{behavior}}(f, X) + w_s \text{err}_{\text{social}}(V, G), \quad (5)$$

where w_s is a regularization hyperparameter. In the following subsection we will discuss choices for $\text{err}_{\text{behavior}}$, as well as optimization techniques used to find the parameters in practice.

Similarity-based error measures related to Equation 3 have been proposed by others, typically based not on a social graph but instead on measured similarities between users. For example, the measured Pearson correlation of item ratings can be used as a similarity measure $\text{Sim}(i, j)$ between users i and j , and this can be incorporated using a measure such as [12]:

$$\sum_{i=1}^{|\mathcal{U}|} \left\| V_i - \frac{\sum_{j,G(i,j)=1} \text{Sim}(i, j) \times V_j}{\sum_{j,G(i,j)=1} \text{Sim}(i, j)} \right\|_2^2. \quad (6)$$

However, accurately estimating similarities from data is often unreliable due to sparsity, especially in the CR setting. Moreover, such measures make it difficult to easily recommend items to newer users; without a long history of ratings, we cannot know which established users they are similar to. On the other hand, SCR requires an external source of information in the form of a social graph. But social networks are increasingly ubiquitous, and, since they are by nature centralized, can often be reliable even when extensive training data for a specific CR task is not yet available.

SCR can be viewed as a blend of social networking, collaborative filtering, and information retrieval. As a side benefit, in addition to providing improved recommendations for users under particular query contexts, SCR can potentially be used in the inverse to recommend new social links between users with similar preferences. In this way SCR can strengthen the social network and improve its own predictions in the future.

2.3 Learning

The goal of SCR learning is to (efficiently) find parameters S , V , T , and U_u that minimize the objective in Equation 5. In this section we describe the formal learning setup, the specific behavioral measures used in our experiments, and the algorithm used to optimize the model parameters.

We assume we are given a training set X containing N training examples:

$$X = \{(q_i, u_i, a_i, w_i)\}_{i=1,2,\dots,N}, \quad (7)$$

where $q_i \in \mathcal{Q}$ is a query, $u_i \in \mathcal{U}$ is a user, $a_i \in \mathcal{A}$ is an item, and $w_i \in \mathbb{R}_{>0}$ is a measure of relevance for the item a_i given the user u_i and the query q_i . Importantly, we assume that the weights w_i always have a positive connotation; that is, triples (q, u, a) that do not appear in the training set implicitly have a weight of zero, and are therefore dispreferred to triples that do appear. For instance, in our experiments, w_i will be derived from the number of times a user listens to a particular musical artist.

The behavioral part of the objective, which measures the compatibility of the scoring function f (defined by the model parameters) with the training set X , can take a variety of forms depending on the setting. As noted earlier, we will focus on top- k ranking losses that optimize the most important aspects of the model, rather than, say, filling out all entries of the tensor \mathbf{R} .

Following Weston et al. [24] we define the vector $\bar{f}(q, u)$, which contains predictions for all items in the database given query q and user u . The a^{th} entry of $\bar{f}(q, u)$, denoted $\bar{f}_a(q, u)$, is equal to $f(q, u, a)$.

With this notation, we can define the *Weighted Approximate-Rank Pairwise (WARP) Loss*, introduced in [23]:

$$\text{err}_{\text{WARP}}(f, X) = \sum_{i=1}^N L(\text{rank}_{a_i}(\bar{f}(q_i, u_i))) . \quad (8)$$

Here $\text{rank}_{a_i}(\bar{f}(q_i, u_i))$ is the margin-based rank of item a_i ,

$$\text{rank}_{a_i}(\bar{f}(q_i, u_i)) = \sum_{b \neq a_i} \mathbb{I}[1 + \bar{f}_b(q_i, u_i) \geq \bar{f}_{a_i}(q_i, u_i)] , \quad (9)$$

where $\mathbb{I}[\cdot]$ is the indicator function, and L is a loss function:

$$L(k) = \sum_{i=1}^k \alpha_i \quad (10)$$

$$\alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \dots \geq 0 , \quad (11)$$

with the values of α_r determining the additional penalty for each successive reduction in rank. We choose $\alpha_r = 1/r$, which gives a smooth weighting over positions while assigning large weights to top positions and rapidly decaying weights to lower positions.

Intuitively, the WARP loss prefers that the item a_i is always ranked highest. For each training example $i = 1, \dots, N$, the positive item a_i is compared pairwise with all other (negative) items in the database. If the score of another item is less than a margin of one from the score of a_i , this pair incurs a cost. The WARP loss determines this cost based on the corresponding items' ranking positions and the choice of α parameters.

We use the WARP loss in our experiments for comparison with prior work. However, in our setting it ignores the relevance scores w_i that are part of the training set; this can be inefficient, since the optimization cannot focus on the most important training examples. We thus propose a modified behavioral measure that we refer to as the weighted WARP loss:

$$\text{err}_{\text{weighted}}(f) = \sum_{i=1}^N w_i L(\text{rank}_{a_i}(\bar{f}(q_i, u_i))) , \quad (12)$$

where rank and L are defined as before. In our results, we refer to models learned under this loss as SCR-weighted, while models trained under the standard WARP loss are referred to simply as SCR. We will derive the optimization procedure for the weighted WARP loss, since the standard WARP loss is a special case.

To minimize Equation 5, we employ stochastic gradient descent, choosing at each iteration a single training instance i uniformly at random from the training set. We then seek to optimize the objective for this single example; that is, we minimize

$$w_i L(\text{rank}_{a_i}(\bar{f}(q_i, u_i))) + w_s \sum_{v, G(u_i, v)=1} \|1 - \sigma(V_{u_i}^\top V_v)\|^2 . \quad (13)$$

Because it is expensive to compute the exact rank of an item a_i when the total number of items is very large, the optimization procedure includes a sampling process at each

step, as introduced in [23]. For the training sample i chosen at the current iteration, negative items b are sampled uniformly at random from \mathcal{A} until a pairwise violation is found—that is, until $1 + f(q_i, u_i, b) > f(q_i, u_i, a_i)$. If K steps are required to find such a b , then the rank of a_i can be approximated as

$$\text{rank}_{a_i}(\bar{f}(q_i, u_i)) \approx \left\lfloor \frac{|\mathcal{A}| - 1}{K} \right\rfloor , \quad (14)$$

where $\lfloor \cdot \rfloor$ is the floor function.

Following Weston et al. [23], the single-instance objective becomes

$$w_i L\left(\left\lfloor \frac{|\mathcal{A}| - 1}{K} \right\rfloor\right) \cdot |1 - f(q_i, u_i, a_i) + f(q_i, u_i, b)| + w_s \sum_{v, G(u_i, v)=1} \|1 - \sigma(V_{u_i}^\top V_v)\|^2 . \quad (15)$$

Rewriting Equation 15, we have

$$C_i(1 + (S_{q_i}^\top U_{u_i} + V_{u_i}^\top)(T_b - T_{a_i})) + w_s \sum_{v, G(u_i, v)=1} \|1 - \sigma(V_{u_i}^\top V_v)\|^2 , \quad (16)$$

where $C_i = w_i \cdot L(\lfloor \frac{|\mathcal{A}| - 1}{K} \rfloor)$. To speed up each gradient step, we only update the variables associated with the current violation pair; that is, we only update S_{q_i} , V_{u_i} , T_{a_i} , T_b , and U_{u_i} . (In particular, we do not update the representations u_i 's friends V_v for $G(u_i, v) = 1$.) Now we can simply take the gradient of (16) to perform an update.

The update for user u_i 's low-dimensional embedding is

$$V_{u_i} \leftarrow V_{u_i} - \eta \left(C_i(T_b - T_{a_i}) + w_s \sum_{v, G(u_i, v)=1} (-2c\sigma(V_{u_i}^\top V_v) (1 - \sigma(V_{u_i}^\top V_v))^2) \cdot V_j \right) , \quad (17)$$

or equivalently

$$V_{u_i} \leftarrow V_{u_i} - \eta C_i(T_b - T_{a_i}) + w'_s \sum_{v, G(u_i, v)=1} b_v \cdot V_v , \quad (18)$$

where $b_v = 2c\sigma(V_{u_i}^\top V_v)(1 - \sigma(V_{u_i}^\top V_v))^2 > 0$, $w'_s = \eta w_s$, and η is a learning rate parameter. (Recall that c is a hyperparameter for the sigmoid function.) Thus at each gradient step, the user's low-dimensional embedding is updated toward the weighted mean of his or her friends' embeddings.

Similarly, we can derive updates for the remaining parameters:

$$S_{q_i} \leftarrow S_{q_i} - \eta \left(C_i(U_{u_i}(T_b - T_{a_i})) \right) \quad (19)$$

$$T_{a_i} \leftarrow T_{a_i} - \eta \left(C_i(-U_{u_i}^\top S_{q_i} - V_{u_i}) \right) \quad (20)$$

$$T_b \leftarrow T_b - \eta \left(C_i(U_{u_i}^\top S_{q_i} + V_{u_i}) \right) \quad (21)$$

$$U_{u_i} \leftarrow U_{u_i} - \eta \left(C_i \left(S_{q_i}(T_b - T_{a_i})^\top \right) \right) . \quad (22)$$

Finally, we constrain the parameters using

$$\|S_i\|_2 \leq C_S, \quad i \in \{1, \dots, |Q|\} \quad (23)$$

$$\|T_i\|_2 \leq C_T, \quad i \in \{1, \dots, |A|\} \quad (24)$$

$$\|V_i\|_2 \leq C_V, \quad i \in \{1, \dots, |\mathcal{U}|\} \quad (25)$$

and project the parameters back on to the constraints at each step.

3. SOCIAL DATA ANALYSIS

Before showing results that compare our proposed approach to existing state-of-the-art methods, we first experimentally validate the fundamental assumption that friends, on average, have more in common than non-friends.

3.1 Last.fm dataset

In our experiments we use a real-word music dataset obtained from the Last.fm music website, *hetrec2011-lastfm-2k* [4]. In this dataset each user is described by his or her listen counts for all musical artists in the database (items, in our setting), as well as any artist tags that might have been provided by each user.

While the data contain more than ten thousand unique tags across all users, the vast majority of tags are used by only one user. Typically these tags appear to be personal "notes" rather than widely used genre distinctions. To remove this noisy information, we throw out tags that are less frequent, keeping only the top 30 most common tags. These tags were all used by at least 165 unique users, and generally correspond to genres of music; for example, the top 5 most popular tags are "rock", "pop", "alternative", "electronic" and "indie". The Last.fm dataset contains listening histories for 1892 users and 17632 music artists. A social graph is also included; on average each user has 13.44 friends.

3.2 Shared musical interests

Do friends share more preferences than non-friends? This is a key question for our approach. If the answer is no, it may not be useful to include social networks as a predictor variable in recommendation systems. To estimate the similarity between two users' tastes for music, we compute the *listened artists overlap ratio*, defined as

$$Sim(i, j) = \frac{|A_i \cap A_j|}{|A_i \cup A_j|} \in [0, 1] \quad \forall i, j, \quad (26)$$

where A_i is the set of artists listened to by user i .

The overlap ratios are computed for all $\binom{|\mathcal{U}|}{2}$ user pairs. We then divide the range $[0, 1]$ of possible ratios evenly into 100 intervals, and calculate the fraction of the user pairs falling in each interval that are linked in the social graph. Intuitively, we hope that users with greater similarity are more likely to be friends. The result is shown in Figure 1 (a). The percentage of realized *friend* relations increases sharply as the artist overlap ratio increases.

To reinforce this analysis, we also compute the average similarity between each user i and his or her friends, as well as the average similarity between user i and all other non-friend users, denoting the two numbers as β_{friend}^i and $\beta_{non-friend}^i$. Figure 1 (b) shows the values of β_{friend}^i and $\beta_{non-friend}^i$ averaged over the set of users with a given number of friends; this number varies on the horizontal axis. We can see that, no matter how many friends a user has, on average the user has more in common with his or her friends

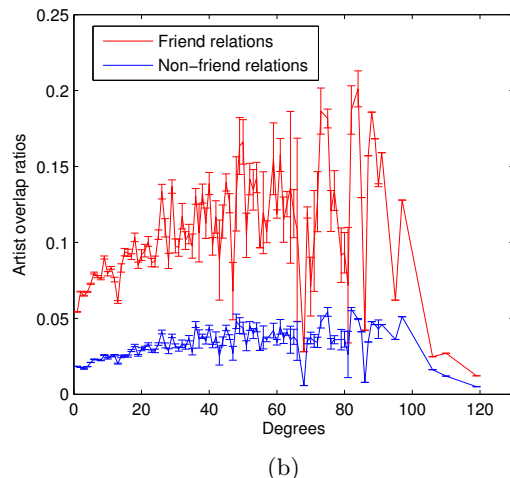




Figure 1: (a) For user pairs having listened artist overlap ratios within the same interval, the proportion of friend relations among these pairs is shown. (b) Average artist overlap ratios among friends (red) and non-friends (blue) for users having different numbers of friends (degree).

than with non-friends; moreover, the size of this effect increases for users with more friends. Overall, these analyses support our assumptions regarding the use of social networks for recommendation and retrieval tasks on this dataset.


4. EXPERIMENTS

We next compare the SCR approach with other state-of-the-art algorithms. For the Last.fm dataset described in the previous section, a *query* \times *user* \times *item* triple corresponds to a *genre* \times *user* \times *artist*, where genres are obtained from the set of filtered user tags.

We preprocess the data set in two ways to obtain listening counts for each such triple. First, if a user u has listened to an artist a and assigned multiple genres, for example *rock* and *indie*, then u 's listening counts for a are evenly distributed to triples $(rock, u, a)$ and $(indie, u, a)$. If the user has not assigned any genre to an artist, the genres of a are



(a)




(b)

Figure 2: (a) Recall@30 of SCR-weighted using different embedding dimensions. (b) Weighted Recall@30 of SCR-weighted using different embedding dimensions.


those assigned to a by other users, and the listening count of a is distributed to each triple according to how frequently the genre appears. If no user has ever assigned any genre to a , the genres of a are defined as the genres used by u for any artist, and the listening count is again prorated to the triples.

Second, since we are interested in ranking artists given a particular user and query, we normalize listening counts of triples having the same u and q so that their weights sum to 1. In the end we have 389,405 data points of the form (q, u, a, w) , where w is the normalized listen count of artist a by user u in genre q .

Since our main goal is to show how social information can help compensate for data sparsity in a collaborative retrieval task, we identify a series of subsets of the Last.fm data that correspond to increasingly less compact social networks. We use a standard implementation of hierarchical clustering on the complete social adjacency matrix to select subsets of users that exhibit significant internal social structure; the number of users in these sets varies from 200 to 1000 (see Table 1). For each user set, the corresponding set of items contains all artists listened to by one or more of the selected users. In this way, the number of artists grows organically



(a)



(b)


Figure 3: (a) Recall at different values of k . (b) Weighted recall at different values of k .

with the number of users. As in Section 3, we use the 30 most frequent genre tags as our query set.

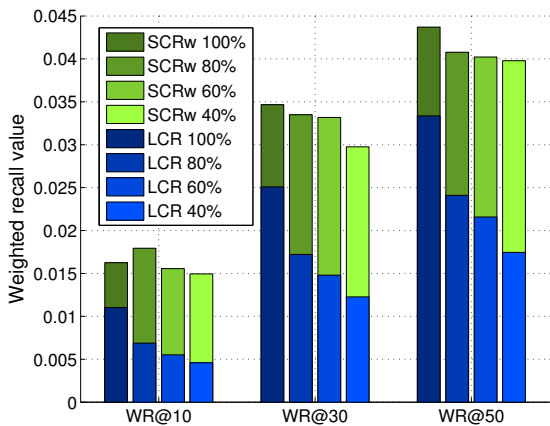
The resulting datasets are referred to as *Compact-lastfm-N*, where N denotes the number of users in the dataset. Their statistics are shown in Table 1. By construction, users in the smaller datasets have higher average numbers of within-set friends. This means that the smaller sets are more tightly connected, which may make them more amenable to social regularization. Conversely, the larger datasets are sparser and may be more representative of large-scale social networks. Note that the *density* of social links falls with the number of users even if the average number of within-set friends stays constant, thus the largest datasets are in fact quite a bit more sparse than the smallest ones. We will show how the relative performance of SCR changes as these qualities are varied.

4.1 Evaluation

To evaluate the performance of each algorithm, for a given test sample (q, u, a, w) we first compute $f(q, u, i)$ for $i =$



(a)




(b)

Figure 4: (a) Recall for training data of different reduced sizes. (b) Weighted Recall for training data of different reduced sizes.


$1, \dots, |\mathcal{A}|$ and sort the artists in descending order of score. We then measure $\text{recall}@k$, which is 1 if artist a appears in the top k , and 0 otherwise, and report the mean $\text{recall}@k$ over the whole test dataset. As a secondary measure we report *weighted* $\text{recall}@k$, which is w if artist a appears in the top k , and 0 otherwise. Mean weighted $\text{recall}@k$ thus not only measures how many triples are ranked in the top k , but the quality of these test triples.

4.2 Results

We compare SCR with state-of-the-art algorithms used for collaborative retrieval as well as traditional collaborative filtering. Popular matrix factorization methods such as SVD and NMF are often used for collaborative filtering; these methods optimize the deviation of the rating matrix from entries supplied in the training set. However, standard SVD and NMF techniques are not directly applicable to tensors. Instead, we perform NMF on the $|\mathcal{Q}|$ different user \times artist matrices to compute the rank of a among all artists given



(a)



(b)

Figure 5: (a) Recall@30 for datasets of different sizes. (b) Weighted Recall@30 for datasets of different sizes.

q and u . We also compare to latent collaborative retrieval (LCR).

The dimension of the embeddings for all methods is chosen to be 30; as shown in Figure 2, this choice yields approximately optimal performance for SCR-weighted; however, the results are not qualitatively different for other choices of embedding dimension. The hyperparameters w_s , η , and c , along with constraint parameters C_S , C_S and C_V , are chosen separately for each method (as applicable) using the validation set. Since matrix factorization approaches are not specially designed for tensors and typically show worse performance than LCR [24], we only present results for NMF, which performed the best. We use the NMF implementation from <http://www.csie.ntu.edu.tw/~cjlin/nmf/>. For each experiment, 60% of the samples are used for training (or less; see below), 20% are used for validation, and 20% are used for testing.

We begin with results for the smallest dataset, *Compactlastfm-200*, which is small enough to be practical for all methods. The resulting $\text{recall}@k$ and weighted $\text{recall}@k$ for different k are shown in Figure 3. SCR (weighted or unweighted) outperforms the baselines on this top k ranking problem; note that SCR-weighted outperforms SCR under

Dataset	users	items (artists)	queries (tags)	samples	Average # of friends
<i>lastfm-2k</i>	1892	17632	11946	186479	13.44
Compact- <i>lastfm-200</i>	200	2392	30	29850	28.54
Compact- <i>lastfm-300</i>	300	3299	30	44318	29.79
Compact- <i>lastfm-400</i>	400	4091	30	58098	29.10
Compact- <i>lastfm-500</i>	500	4928	30	72125	27.81
Compact- <i>lastfm-600</i>	600	5765	30	85522	26.32
Compact- <i>lastfm-700</i>	700	6454	30	98367	24.90
Compact- <i>lastfm-800</i>	800	7071	30	111062	23.57
Compact- <i>lastfm-900</i>	900	7782	30	124005	22.23
Compact- <i>lastfm-1000</i>	1000	8431	30	137518	21.01

Table 1: Last.fm dataset statistics




Figure 6: Runtime required for each training iteration of LCR and two versions of SCR on the Compact-*lastfm-200* dataset.

the weighted recall criterion, which makes sense since it incorporates relevance scores in the loss function.

Since we expect social information to be particularly useful when data are sparse, we also show results of recall@ k and weighted recall@ k for different amounts of training data (100%, 80%, 60% and 40% of the total training data) in Figure 4. Notice that the performance gap between SCR-weighted and LCR becomes larger as the number of available training examples is reduced, suggesting that our proposed algorithm can be especially useful for predicting the interests of new users or infrequent users when social network information is available.

Moving to the larger datasets, computation time increasingly becomes an issue for the matrix factorization approaches, so we focus on only two algorithms: SCR-weighted and LCR. Figure 5 shows recall results for all of the compact datasets. Note that the performance gap between SCR and LCR narrows slightly but remains significant even as the size of system becomes larger and the density of social links decreases. It may be counterintuitive that performance decreases (at least for unweighted recall) as the size of the dataset grows; however, since the number of artists grows with the number of users, the prediction problem is becoming more difficult

at the same time. These results suggest that, while a dense social network may improve the relative performance of SCR, it retains significant advantages even in larger, sparser settings.

Finally, we show in Figure 6 the runtimes for each stochastic gradient training iteration of SCR and LCR on the Compact-*lastfm-200* dataset; SCR is dramatically faster, despite using essentially similar optimization techniques. This is because the runtime is dominated by the sampling procedure used to estimate the rank function. LCR promotes the observed items to high positions quickly, thus making subsequent iterations quite slow. SCR, on the other hand, has additional regularization that appears to prevent this situation. Combined with the performance improvements discussed above, this is a significant practical advantage.

5. CONCLUSIONS AND FUTURE WORK

We proposed a new model that blends social networking, information retrieval, and collaborative filtering. Our SCR model uses social networks to improve performance on collaborative retrieval problems, which we believe are rapidly becoming important in practice. The proposed method outperforms state-of-the-art CR and CF approaches, and shows that users tend to share interests with friends. We believe the combination of recommendation and retrieval using social information will become increasingly useful; going forward we hope to develop a two-pass version of the SCR algorithm that helps predict interest commonalities between friends, and can be used to prune out edges on the social graph that may work against achieving good performance.

References

- [1] Punam Bedi, Harmeet Kaur, and Sudeep Marwaha. Trust based recommender system for semantic web. In *proceedings of the 2007 International Joint Conferences on Artificial Intelligence*, pages 2677–2682, 2007.
- [2] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.
- [3] John S Breese, David Heckerman, and Carl Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth conference on Uncertainty in artificial intelligence*, pages 43–52. Morgan Kaufmann Publishers Inc., 1998.

- [4] Ivan Cantador, Peter Brusilovsky, and Tsvi Kuflik. Second workshop on information heterogeneity and fusion in recommender systems (hetrec2011). In *Proceedings of the fifth ACM conference on Recommender systems*, pages 387–388. ACM, 2011.
- [5] David Carmel, Naama Zwerdling, Ido Guy, Shila Ofek-Koifman, Nadav Har’El, Inbal Ronen, Erel Uziel, Sivan Yogev, and Sergey Chernov. Personalized social search based on the user’s social network. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 1227–1236. ACM, 2009.
- [6] Alexandros Karatzoglou, Xavier Amatriain, Linas Baltrunas, and Nuria Oliver. Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 79–86. ACM, 2010.
- [7] Henry Kautz, Bart Selman, and Mehul Shah. Referral web: combining social networks and collaborative filtering. *Communications of the ACM*, 40(3):63–65, 1997.
- [8] Ioannis Konstas, Vassilios Stathopoulos, and Joemon M Jose. On social networks and collaborative recommendation. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 195–202. ACM, 2009.
- [9] Greg Linden, Brent Smith, and Jeremy York. Amazon.com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE*, 7(1):76–80, 2003.
- [10] Hao Ma, Haixuan Yang, Michael R Lyu, and Irwin King. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings of the 17th ACM conference on Information and knowledge management*, pages 931–940. ACM, 2008.
- [11] Hao Ma, Irwin King, and Michael R Lyu. Learning to recommend with social trust ensemble. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, pages 203–210. ACM, 2009.
- [12] Hao Ma, Dengyong Zhou, Chao Liu, Michael R Lyu, and Irwin King. Recommender systems with social regularization. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 287–296. ACM, 2011.
- [13] Paolo Massa and Paolo Avesani. Trust-aware collaborative filtering for recommender systems. In *On the Move to Meaningful Internet Systems 2004: CoopIS, DOA, and ODBASE*, pages 492–508. Springer, 2004.
- [14] Paolo Massa and Paolo Avesani. Trust-aware recommender systems. In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 17–24. ACM, 2007.
- [15] John O’Donovan and Barry Smyth. Trust in recommender systems. In *Proceedings of the 10th international conference on Intelligent user interfaces*, pages 167–174. ACM, 2005.
- [16] Sanjay Purushotham, Yan Liu, and C-C Jay Kuo. Collaborative topic regression with social matrix factorization for recommendation systems. *arXiv preprint arXiv:1206.4684*, 2012.
- [17] Ruslan Salakhutdinov and Andriy Mnih. Probabilistic matrix factorization. *Advances in neural information processing systems*, 20:1257–1264, 2008.
- [18] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, pages 285–295. ACM, 2001.
- [19] D Seung and L Lee. Algorithms for non-negative matrix factorization. *Advances in neural information processing systems*, 13:556–562, 2001.
- [20] Chirag Shah. Lessons and challenges for collaborative information seeking (cis) systems developers. *Group*, 2009.
- [21] Xiaoyuan Su and Taghi M Khoshgoftaar. A survey of collaborative filtering techniques. *Advances in Artificial Intelligence*, 2009:4, 2009.
- [22] Panagiotis Symeonidis, Alexandros Nanopoulos, and Yannis Manolopoulos. Tag recommendations based on tensor dimensionality reduction. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 43–50. ACM, 2008.
- [23] Jason Weston, Samy Bengio, and Nicolas Usunier. Large scale image annotation: learning to rank with joint word-image embeddings. *Machine learning*, 81(1):21–35, 2010.
- [24] Jason Weston, Chong Wang, Ron Weiss, and Adam Berenzweig. Latent collaborative retrieval. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 9–16, 2012.
- [25] Liang Xiong, Xi Chen, Tzu-Kuo Huang, Jeff Schneider, and Jaime G Carbonell. Temporal collaborative filtering with bayesian probabilistic tensor factorization. In *Proceedings of SIAM Data Mining*, volume 2010, 2010.
- [26] Vincent W Zheng, Bin Cao, Yu Zheng, Xing Xie, and Qiang Yang. Collaborative filtering meets mobile recommendation: A user-centered approach. In *Proceedings of the 24rd AAAI Conference on Artificial Intelligence*, 2010.