Personalized Recommendation of Research Papers by Fusing Recommendations from Explicit and Implicit Social Networks

Shaikhah Alotaibi
University of Saskatchewan
181 Thorvaldson Building
110 Science Place
Saskatoon, SK, Canada
Shaikhah.otaibi@usask.ca

Julita Vassileva
University of Saskatchewan
178 Thorvaldson Building
110 Science Place
Saskatoon, SK, Canada
jiv@cs.usask.ca

ABSTRACT
Combining social network information with collaborative filtering recommendation algorithms has helped to alleviate some drawbacks of collaborative filtering, for example, the cold start problem, and has increased the accuracy of recommendations. However, the user coverage of recommendation for social-based recommendation is low as there is often insufficient data about explicit social relationships among users. In this paper, we fuse recommendation that uses explicit social relations (friends and co-authors) with recommendations that use implicit social relations aiming to increase the user coverage with minimum recommendation accuracy loss. We found that fusing recommendations from friends with recommendations using implicit social networks increases both accuracy and user recommendation coverage while fusing recommendation from co-authors increase the coverage.

CCS Concepts
• Information system applications → Collaborative and social computing, Data mining • World Wide Web → Web searching and information discovery, Web applications

Keywords
Social network; implicit social network; hybrid recommendation; paper recommendation; social bookmarking websites; collaborative filtering.

1. INTRODUCTION
Scholarly papers both help to update researchers on new research in their areas of interest and serve as a directory of other researchers with similar interests with whom researchers can collaborate. However, as publishers, online journals, and conferences proliferate, the number of new published papers has become overwhelming. For this reason, many recommender systems (RSs) have been proposed to help readers by suggesting a list of potential papers of interest. The two main algorithms used by RSs are content-based filtering (CBF) and collaborative filtering (CF). CBF is based on information retrieval techniques that compare a paper's features (e.g., title, abstract, keywords, publication year) with the researchers' features (e.g., interests or previous search queries) to find matches [2]. In contrast, CF (e.g., [14]) uses the similarity of paper ratings to find users similar to the target user and recommend papers that these users have liked. Hybrid recommending approaches (e.g., [25]) use a combination of the CBF and CF approaches to alleviate the drawbacks of both approaches.

Another way to overcome one or more of the CF drawbacks (e.g. cold start problem or data sparsity) is to exploit the social ties between users in the recommendations. With the advent of social networks in applications such as social bookmarking systems (e.g., CiteULike, Mendeley), which researchers often use to manage their digital paper repositories and bookmark libraries, users can be connected through different social relations. By knowing that two users are connected, one can infer that they possibly share interests and therefore recommend items from other connected users. A social bookmarking service provides many clues for interest similarities between users based on their behavior in the system and their publication authorship. Surprisingly, however, none of the popular social bookmarking tools have used the wealth of social data they store to build a social RS. However, there are some studies that incorporate social information into CF techniques to increase the recommendation accuracy. Although such social recommenders perform well, the social information about users that they require is not readily available for all users. Thus, these social recommenders have lower user coverage [18]. User coverage is the ratio of users who receive nonempty recommended sets to the total number of users [3]. Previous studies also showed that there is a tradeoff between the recommendation accuracy and the coverage of the recommendation [10]. For this reason, in this paper we investigate the user utility expressed in recommendation *accuracy* and *coverage* when different types of social relations are used. We hypothesize that using social data from two social networks where the first has higher recommendation accuracy and the other has higher recommendation coverage will result in recommendation that achieve a balance between both benefits. Most of the previous studies fused different recommendation algorithms using the same data source. However, in our study, we fuse different recommendations that use different types of social relations: explicit social relations which are the relations initiated by one or both users involved in the social relation, and implicit social relations which are computed by machine using data about the user behavior in the social bookmarking tools (e.g. co-bookmarking the same paper). The social networks that we use in this study are: two explicit social networks based on co-authorship and friendship, and three different implicit social networks based on readership, co-readership and tag-based social network.

Section 2 discusses related work. Section 3 describes the social networks used in this paper. Section 4 describes the recommender approaches used. Section 5 explains the experiments, dataset and result analysis. Finally, section 6 discusses our conclusions and future work.

2. RELATED WORK
Although the first social recommendation approach appeared as early as 1997 [11], no agreed-upon definition for social recommendation existed until 2013, when Tang et al. defined social recommendation to be any recommender system that includes social relations as an extra input [24]. Thus, social recommenders are hybrid recommender systems that combine social relationships (e.g. membership,

friendship, following relations, trust relations) with another recommendation method, most commonly CF. Rather than using only the user–item matrix as in traditional CF, a social recommendation mechanism uses two matrices: a user–item matrix, which represents the items that are rated by the user, and a user–user matrix, which represents the social relations between users. Many studies demonstrate that using social information in the recommendation process reduces the effect of the data sparsity and cold start problems [13] and enhances prediction accuracy [20].

There are many approaches combining CF recommender with a social network based on explicit social relations between users (e.g. [3,8,20,26]) for example, following on Twitter or Instagram, being friends on Facebook, or in general connection that is made with the awareness or agreement of both users. For instance, Liu and Lee [20] compared four algorithms to recommend skin products: nearest neighborhood CF, social CF, a combination of nearest neighborhood CF and social CF, and nearest neighborhood CF with an amplification of data from social friends. Groh and Ehmig [8] considered the user's friends to form the user's neighborhood to recommend local clubs using social CF. Yuan et al. [26] tested the effect of two explicit social networks, membership and friendship, fused with conventional CF recommendation methods to recommend music. Bellogin, Cantador and Castells [3] tested different recommendation approaches to recommend music items using tags and social network information.

All previous studies in social recommendations using explicit relations between users are in the taste domains (i.e. recommending music, movies or clubs or restaurants. However, it is difficult to use the recommendation algorithms developed for a taste domain to recommend research papers because in taste domains, the number of ratings for each item is larger than the number of ratings received per research paper. Generally, researchers seem reluctant to rate research papers in bookmarking systems, and there is a lack of explicit ratings in the domain of research paper recommendation. Thus, most of the research done in this area is based on citation networks and implicit feedback about the papers, generated from user actions such as tagging, downloading, or bookmarking.

Existing research has explored also the use of implicit networks in social recommender systems. Implicit social networks are constructed by inferring relationships between users that may not exist in the real world, and the users may be unaware of them. For example, the users that belong to the same neighborhood in a CF could be considered as part of an implicit network constructed by relating uses who gave similar ratings to the same items. These implicit relationships have been often called "trust" [1, 7, 13, 21]. For example, in [1], a trust-aware RS is proposed that uses trust metrics to personalize the recommendations for secure skiing routes by showing information from only users the target user trusts. The trust in Moleskiing is used to alleviate the data sparsity problem using trust propagation to infer the trust values for unknown users. The FilmTrust social Web site system proposed by Golbeck [7] recommends movies using the trust developed between users based on similar movie ratings. A study done by O'Donovan and Smyth [21] incorporates implicit trust values inferred from user ratings into standard CF. Massa and Avesani [13] propose a trust graph–based RS that uses trust values given by users in addition to similarity measures to reduce the data sparseness that affects new users. The results of their experiments, performed on the Epinions dataset, show that trust-aware RSs outperform CF.

Some studies compared the recommendations produced by explicit social networks with those produced by implicit social networks. For example, Guy et al. [9] compared recommendation produced by data from users' friends with recommendation produced using implicit social relations among users based on their behaviors, such as using the same tag or co-bookmarking the same webpage. Then, they compared the results with the recommendation from people who are familiar and similar to other users (i.e., a combination of both previous social networks). They showed that the recommendation from users' friends outperformed the recommendation of the implicit social network. They explained their result with the fact that the recommendations are explained to the users, who can see the picture of the contact who sent the recommendation. However, this result could be explained also with the fact that all users belong to the same community, which, in this case, was the IBM Corporation, so they all have similar interests and also know each other.

A wise social network recommender system (WSNRS) was proposed by Mican et al. [15]. It considered explicit and implicit social relations (e.g., implicit relations based on number of clicks to see other user's profile). First, the algorithm considered the user's connections made up of users who have explicit social relations with the target user. It then considered the interactions between the target user and other users as well as the interactions between the target user and the webpages to calculate a trust value. If the trust value was above average, the target user is an implicit follower of the other user. The recently published resources and the favorably rated resources from the target user's connections are then recommended to the target user. Mican et al. explained and demonstrated this using a case study that was neither evaluated by any evaluation metrics nor compared with any baseline recommendation methods. Thus, the effectiveness of the proposed method is unclear.

Generally, very few studies incorporate social relations in the domain of research paper recommendations. One of the existing works, PubRec, is an RS that suggests to the target user, for a particular paper of interest, the most related papers from the libraries of other users to whom that user is socially connected [22]. PReSA [23] takes advantage of the available data on social bookmarking websites (e.g., CiteULike), such as bookmarked papers, metadata, and users' connections, to recommend papers from the users' connections' libraries that are similar and popular among the users' social connections. Both PubRec and PReSA consider the explicit relationships among users in the recommendation process. Lee and Brusilovsky studied three explicit social networks—watching networks [16], group membership [17], and collaboration networks [18]—to find the extent of interest similarities between users involved in those networks and compare the recommendations watching networks produced to the recommendations traditional CF produced [16]. Their results showed that the watching network cannot compete with CF, that the similarities between users' libraries in group membership networks are insignificant [17], and that the similarity between two users connected using co-authorship networks is comparable to user connections using explicit networks, which require agreement between the parties [18].

All the studies that have been done in the area of exploiting social relations in recommending research papers are based on explicit social networks, which have low coverage. In the next section, we propose three social networks where the social relations between users are inferred based on their publications and their bookmarked papers. We test the recommendation accuracy using these social networks as sources of information and we also test the user coverage. Then, we compare the results with hybrid approaches that fuse recommendation from implicit and explicit social networks to provide a good balance between recommendation accuracy and user coverage.

3. EVALUATED SOCIAL NETWORKS

Three implicit social networks (ISNs) based on users' bookmarking behavior are proposed as candidates to use for recommending research papers. In addition, two commonly used explicit social networks are introduced: co-authorship and friendship. In this section we describe all these five social networks.

3.1 Implicit Social Networks

We consider three different social networks that connect users to each other based on their bookmarking and tagging behavior in social bookmarking tool. First, the *readership ISN* connects users to the authors of the papers that they have bookmarked. We assume that if users bookmark specific papers, interest overlap exists between the bookmarkers and the authors of the papers; this overlap increases with the increase in the number of papers users bookmark from the same author. The relation could be *unidirectional* or reciprocal. The relation is unidirectional if only one of the users in this relation has bookmarked the other user's publications. The relation is *reciprocal* if both users have bookmarked each other's publications. Figure 1 shows the relations in this network, which are depicted as black arrows. For example, the relation between user 3 and user 5 is reciprocal, while the relation between user 3 and user 1 is unidirectional; user 3 is the paper's bookmarker and user 1 is the paper's author. The numbers on the arrows represent the strength of the relations. For example, the strength of the relation between user 3 and user 1 is five, which means that user 3's library contains five bookmarked papers authored by user 1.

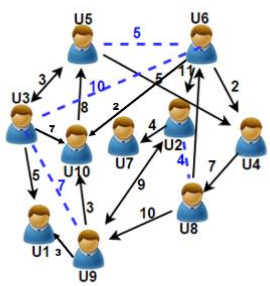

Figure 1. Sample of relations in implicit networks

Second, the *co-readership ISN* connects users who bookmark (and presumably read) papers written by the same authors. If user 1 and user 2 have both bookmarked papers written by user 3, then user 1 and user 2 are connected using the co-readership ISN. This network structure is useful for users who do not yet have publications and therefore cannot have relations in network 1. The assumption is that users who bookmark the same paper(s) also have similar interests. The strength of the relationship is measured by the number of authors whose libraries overlap. Figure 1 shows an example of the relationships in this network in blue. For example, user 5 and user 6 are connected because they both bookmarked papers written by the same authors; the number of overlapping author names here is five. We show only a part of the graph, and it includes only one of those five authors (user 4).

Third, the *tag-based ISN* connects users if they use the same tags to annotate their bookmarked papers. However, we do not check whether users use the same tags to annotate the same papers. We consider the tag similarity between the entire tag cloud associated with each user. We assume that the more similar tags the users have, the higher the interest similarity. While the previous two networks are based on the papers' metadata, this network is based on user-generated data. To build this network, the tags used to annotate the papers are aggregated for each user. The data is preprocessed to make the tags comparable. We follow the method described in [19] to preprocess the tags. All tags are preprocessed by converting them to lowercase, removing the stop words, and then using the porter stemmer tool to remove any additional letters added to the root word to eliminate the effect of the word variation (e.g., the word "social" could have different variations, such as "socialize", "socialization" or "socializing"). The relations in this network also have strengths. The strength of the relation between two users is measured by the number of tags they share. The assumption is that the more tags two users share, the stronger the relationship is between them.

3.2 Explicit Social Networks

We consider two social networks where the relations between users are explicitly defined: co-authorship SN and friendship SN. The *co-authorship* relations between two users emerge when they collaborate in writing and publishing a research paper(s). When two users collaborate in publishing papers, this means that they share similar interest and have strong relationship. The *co-authorship* SN also has a strength represented by how many papers the pair of users has co-authored. The other explicit SN is the *friendship* SN, which has undirected relations. The relation emerges between two users when one user invites the second user to add her to her connection (i.e. contact) list, and the second user accepts the invitation.

4. EVALUATED RECOMMENDING APPROACHES

To determine the effectiveness of the three introduced ISNs as sources for recommendations, we compared three existing commonly used social recommendation approaches: social recommender, combined recommender and amplified recommender. These recommendation approaches were applied previously to datasets that have explicit social relations and numeric ratings of items (i.e. rating of items using Likert scale). We applied the same approaches to a dataset that has implicit social relations and unary ratings of items (i.e. existence of the paper in the user's library).

4.1 Social Recommender

The social recommender was proposed by [3]. It simply replaces the anonymous nearest neighbors in the user-based CF with the target user's social friends in the social network. To apply the social recommender to the proposed ISNs, we found the social friends of each user and used the data from those friends in the same way that anonymous peers in CF are used - by picking the top N peers and using their bookmarked papers to find candidate papers to recommend to the user. However, in the social recommender we replaced the similarity between users that is used in the predication of the target user's rating for unseen items with the weighted strength between users U_i and U_j $WStr_{U_{i,j}}$ calculated as:

$$WStr_{U_{i,j}} = \frac{Str_{U_{i,j}}}{TotalStr_{U_i}}$$

Where $Str_{U_{i,j}}$ is the strength of the relation between U_i and U_j and $TotalStr_{U_i}$ is the sum of all strength values between U_i and all of other users who are connected to her.

4.2 Combined Recommender

The combined recommender integrates neighbors from conventional user-based CF and the target user's social friends to construct a new nearest neighborhood set for the target user [3]. We then used the data from users in the new combined neighbors in the recommendation following the same way as in the social recommender.

4.3 Amplified Recommender

The amplified recommender amplifies the social friends' preferences in CF nearest neighbors [20]. First, the nearest neighborhood peers were identified by CF top-N technique. Then, if the user's social friends were also in the top-N neighbors, we used an amplifying approach to give the preferences from those social friends more weight in the recommendation process. The amplifying function that we used is the one used in [20], which is given by:

$$\text{Min}\ (S_{U_iU_j} \times (1 + \frac{N_{U_iU_j}}{N_{all,U_j}})\ ,1)$$

where U_j is the target user, U_i is one of the U_j's social friends, $S_{U_iU_j}$ is the similarity between U_i and U_j which is calculated by CF approach using the papers that are co-bookmarked by both users, $N_{U_iU_j}$ is the number of interaction between the target U_j and the user's social friend U_i, and N_{all,U_j} is the total number of interactions between the target U_j and all of the user's social friends. Because the similarity value cannot be greater than 1, we chose a minimum value between 1 and the amplified value. The interactions between the target user and one of the user's social friends were based on the type of ISN on which we were trying to apply the approach. For example, if we use the co-readership ISN, the number of interactions equals the number of authors that both users have in common (i.e., the number of authors one or more of whose papers both users bookmark).

5. EXPERIMENTS AND DATASET

5.1 Dataset

We collected the data for this study from the CiteULike.org social bookmarking website. This site has been in active use since November 2004; it currently has 8,217,384 bookmarked papers. It allows social features such as connecting users, watching users (similar to following on Twitter), and sharing references. Users of CiteULike can import scientific reference data from other resources such as PubMed and can assign tags to the bookmarked references for future easy access. Using the snowball method, we crawled the CiteULike website, starting with 500 randomly chosen, recently active users whose publications and bookmark data we collected. Then, we branched to collect the users' data for the users who had bookmarked their publications or who had bookmarked the same papers as the initial users. Table 1 shows the descriptive statistics for the dataset collected between December 2014 and February 2015.

Table 1. Descriptive Analysis of the Dataset

Number of users	13,189
Total number of distinct papers	1,043,675
Total number of publications/bookmarks/tags	19,774/1,323,065/ 3,086,565
Average number of publications/bookmarks/tags per user	1.52/98.79/3.81
Number of unidirectional relations in readership ISN/number of users having unidirectional relations	9,248/4,909
Number of reciprocal relations in readership ISN/number of users having reciprocal relations	141/209
Number of relations in co-readership ISN/number of users in this network	260,361/11,484
Number of relations in tag-based ISN/ number of users in this network	223,405/11,283
Co-authorship SN	
Number of co-authors	247
Total/average number of social relations/ per user	167/1.27
Total/average number of the co-authors' publications/per user	4181/16.93
Total/average number of their bookmarks/per user	43134/174.63
Friendship SN	
Number of users who have friends	2375
Total/average number of bookmarks/per friend	360,715/99.15
Total/average number of friends/per user	6171/0.31

5.2 Metrics of Evaluating Different Recommenders

We compute several metrics from the Informational Retrieval field to measure the prediction accuracy of recommenders. Since the data that we have is bookmarking data, which is considered as unary rating (i.e. presence of absence of rating), the best metrics are precision and recall at top N. It is always assumed that the items with higher ranks in the recommended list of items are more important than items with lower ranks. We calculate the precision and recall for three ranks: top two, top five, and top ten. Then we compare the results among these ranks. Precision@N (reported as P@N in Figure 2) and Recall@N (reported as R@N in Figure 2) are calculated with respect to the rank. For example, if Precision@10 is used, we calculate the ratio of the true recommended items to the top 10 recommended items, and the Recall@10 is the ratio of the number of true recommended items in the top 10 recommended items to the test set. In all of our experiments, we used fivefold cross validation approach where 20% of the user's bookmarks are used as testing data and 80% are used as training data. This process is repeated five times, each time with different test and training sets. Then the accuracy of the prediction is calculated.

When N, the number of recommended items, increases, a trade-off between precision and recall metrics is observed. When N increases, the precision value starts to decrease, while the recall starts to increase. To reduce the effect of the change of the precision and recall by increasing the N value, the F1-measure (reported as F1@N in Figure 2) is used to produce evaluation results that are more universally comparable. F1 can be calculated using the following equation:

$$\text{F1@N}=\frac{2.P@N.R@N}{P@N+R@N}$$

5.3 Experiments and Results

In this section, the conducted experiments are described, and the results for each experiment are explained.

5.3.1 Finding the Best Settings for Each ISN

First, we run the different recommenders described in section 4 for each of the implicit social networks (readership ISN, co-readership ISN and tag-based ISN) as well as for the two explicit social networks (co-authorship SN and friendship SN). To find the best settings for each network, different neighborhood sizes (k value) are tested and different ranked lists are produced (top two, top five, and top ten).

We found that the best performing settings for each network with respect to more metrics are:

- Readership ISN (reciprocal relations): social recommender with K=5
- Readership ISN (unidirectional relations): combined recommender K= 20
- Co-readership ISN: amplified recommender, K= 20
- Tag-based ISN: amplified recommender, K= 20
- Co-authorship SN: amplified recommender, K= all relations
- Friendship SN: amplified recommender, K= all relations

Therefore, we used these settings in the next experiment when fusing data from different social networks. With compatible results with the study in [4], small neighborhood size provided the best accuracy results. In addition, as noted for the explicit social networks (friendship and co-authorship), the best results were achieved by using all the of users' social relations. This is because each user has very few social relations; the average number of relations in friendship networks and co-authorship networks are 0.3 and 1.27 respectively.

5.3.2 Hybrid Recommendation of Explicit and Implicit Social Networks

Then, we used a weighted hybrid recommender to combine the results of recommending research papers using data from explicit and implicit social networks. Even though there are many hybrid approaches [6], we prefer to use the weighted hybrid approach because it brings together all the capabilities of the combined approaches in a straightforward and easy to perform way. It is a linear combination that aggregates the prediction score from different recommendation approaches using a different weight for each recommendation approach. The hybrid recommendation is calculated from the linear combination of different recommendations using

$$Wrec_i = \sum_{S_j \in S} (W_{rec_{i,S_J}} \,.\, W_{S_J})$$

where W_{S_J} is the weight for the recommender S_j, and its value ranging from 0.1 to 0.9, and the sum of all weights is equal to 1. Since the optimum weight is usually derived by examining the performance of all possible combinations [6], we used all the combinations from 0.1 to 0.9 by gradually increasing the weight of the first recommender by increments of 0.1. We first tested the hybrid approach of the co-authorship network (explicit) with every implicit social network, then we tested the friendship network (explicit) with every implicit social network. However, we used a modified version of the weighted sum approach called cross-source hybrid [5]. Cross-source hybrid approach favors items that are recommended by both approaches. We agree that items that are recommended by implicit social network recommender and explicit social network recommender are more important than items that are recommended by only one recommender. Therefore, the above equation for weighted sum hybrid approach is modified as follows:

$$Wrec_i = \sum_{S_j \in S} (W_{rec_{i,S_J}} \,.\, W_{S_J}) \,.\, |S_{rec_i}|$$

Where $|S_{rec_i}|$ is the number of recommenders that recommend item i. We use the cross-source hybrid if the user has relations in both social networks. However, we used weight 1 for the recommendation if the user has relations in only one of social networks. For instance, if we aim to fuse the recommendations produced by co-authorship explicit network with recommendation produced by co-readership ISN, but the user has only relations in co-readership, we use the weight 1 for the recommendation produced by co-readership and completely ignore the co-authorship ISN for this specific user. We used this approach to make the recommendations more personalized. The best weight combination for each hybrid approach is shown in Table 2. When recommendations using co-authorship network are fused with recommendations from reciprocal readership ISN, the maximum accuracy is reached when the recommendations from the co-authorship network are given high weight, 0.8. However, when co-authorship network recommendation is fused with other ISNs, the best accuracy achieved when the weight of co-authorship was 0.3 in the case of unidirectional readership ISN and 0.1 in the case of co-readership ISN and tag-based ISN. This is because there is a high overlap between the co-authorship social relations and the reciprocal readership relations; 58.68 percent of the relations in the co-authorship network was discovered by the reciprocal relations in readership network. The effect of the co-readership is less visible in the other networks, and that might be because there is a huge gap between the small number of relations in co-authorship network and the large number of relations in the other networks.

When recommendations from friendship network are fused with recommendations from implicit social networks, we can notice that the maximum accuracy of the recommendations occurs when the weight of the friendship network is higher than the weight of implicit networks.

Table 2: The optimum weights for each hybrid approach (ISN weight, explicit SN weight)

	Co-authorship	Friendship
Readership ISN (Reciprocal)	(0.2,0.8)	(0.3,0.7)
Readership ISN (unidirectional)	(0.7,0.3)	(0.3,0.7)
Co-readership ISN	(0.9,0.1)	(0.1,0.9)
Tag-based ISN	(0.9,0.1)	(0.1,0.9)

The results of the best weight combinations are used in the next experiment described in the next subsection 5.3.3.

5.3.3 Comparison between Recommendations from Different ISNs with Friendship or Co-authorship SNs

We conducted an experiment to compare the recommendation using only ISN data, with the hybrid approach that combine recommendations from ISNs with each of the two explicit networks – the co-authorship network, or the friendship network, respectively. The results shown in Figure 2 reveal that the best prediction accuracy is achieved when the recommendation from the friendship network is fused with recommendation from ISN; this is true for all implicit social networks and for all measures (precision, recall and F1 at top ten). However, the co-readership ISN did not help in increasing the prediction accuracy. In most of the cases, the prediction accuracy stayed the same. In the readership ISN, the prediction accuracy slightly decreased when recommendation from co-authorship SN is fused. The only case that the precision increased is when recommendation from unidirectional readership SN is used with co-authorship SN recommendation.

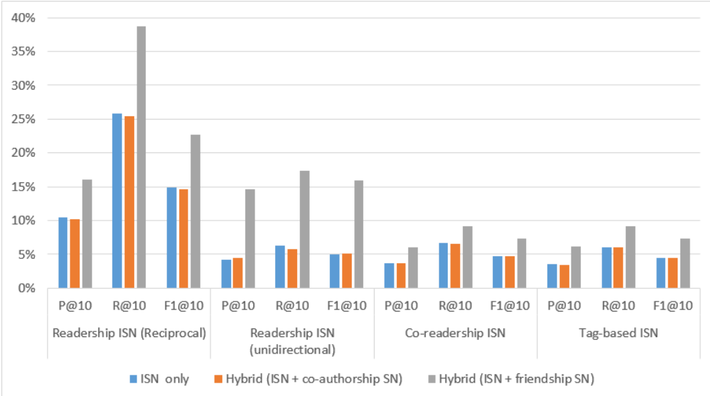

Figure 2: Comparison between using recommendations from ISNs only, or fusing the recommendation with co-authorship SN and friendship network

5.3.4 User Coverage

While measuring the prediction accuracy of recommendation to filter several recommendation approaches is important, it is not the only way to evaluate the performance of a certain recommendation approach. Non-performance measures, such as serendipity, diversity, novelty, or coverage, can also evaluate recommendation approaches [12]. One measure that compares the capability of different recommending approaches to produce recommendations for a larger set of users is the coverage measure, which is the ratio of users who receive nonempty recommended sets to the total number of users. The more coverage provided, the better the recommending algorithm.

We found that the co-readership ISN had the highest user coverage (87.25 percent), higher than the tag-based ISN (85.55 percent), the unidirectional readership ISN (37.22 percent), and the reciprocal readership ISN (1.59 percent). We found that only 18 percent of users have explicit social relations and the average number of social relations per user is only 0.31. The co-authorship explicit SN has a very low coverage (1.873 percent). A tradeoff is noticed between the prediction accuracy and the user coverage: the more accurate the prediction, the smaller the user coverage.

Table 3 shows the coverage of different social networks and compares them to the coverage of the hybrid approaches. The recommendation coverage increases when recommendations from explicit and implicit SNs are combined. However, the maximum coverage is reached when recommendation from the friendship SN is fused with any of the ISNs. This is true for all of ISNs. For example, the increase in the coverage for the reciprocal readership ISN when the recommendation is fused with the friendship SN is more than 16%, while the increase in the coverage when the recommendation fused with the co-authorship SN is only 0.59%.

Fusing recommendation from friendship SN increases both the prediction accuracy and the recommendation coverage. The unidirectional readership ISN is the network that improved the most from fusing recommendation from friendship SN (F1-measure increase of almost 11%), then the reciprocal readership ISN (7.8 %), tag-based ISN (2.9% increase), and finally, the co-readership ISN (with 2.5% increase). Even though, fusing recommendation from co-authorship SN did not improve the recommendation accuracy, it improves the recommendation coverage.

Table 3: Comparison between the user coverage of using different hybrid approaches and using recommendation from ISNs alone

	ISN data only	**Hybrid with co-authorship SN**	**Hybrid with friendship SN**
Reciprocal readership ISN	1.59%	2.18%	18.58%
Unidirectional readership ISN	37.22%	37.43%	45.14%
Co-readership ISN	87.25%	88.9%	92.71%
Tag-based ISN	85.55%	85.62%	86.57%

6. Conclusions and Future Work

We compared the prediction accuracy and the recommendation coverage in three implicit social networks: readership ISN (reciprocal relations or unidirectional relations), co-readership ISN and tag-based ISN. Then we fused recommendation from these networks with two social networks that are based on explicit social relations between users, namely: friendship SN and co-authorship networks. Weighted sum approach was used to fuse the recommendations from two sources (implicit and explicit social networks). The experiments showed that fusing the recommendations from each ISN with recommendations from either the friendship or co-authorship explicit network is beneficial in increasing the user coverage. In addition, the prediction accuracy of all the recommendations from ISNs improved when fused with the friendship explicit SN, but fusing with the co-authorship SN did not help in improving the recommendation accuracy. Therefore, a hybrid approach that fuses recommendations from explicit social networks such as friendship can increase both prediction accuracy and user coverage. This finding is beneficial for recommender algorithm designers to consider hybrid approaches that take into account different social relations of users. Such social relations can be harvested from bookmarking systems that allow synchronizing information from different systems such as (CiteULike.org and Mendeley.com). Our findings are also beneficial for recommender systems interface designers, highlighting the need for allowing users to set weights of different recommendation sources and show the recommended list with explanations.

In the future, we want to generalize our findings by testing the proposed implicit social networks with other datasets and/or with different applications. In this paper, we used the same weights for all users to combine recommendations from different resources, which limits the personalization capabilities. In the future, we want to test using dynamic weights that are based on each user's features such as social network features (e.g. number of incoming/outgoing social relations). We also want to test the recommendations produced by ISNs with real users to test the user perception of and/or satisfaction with recommendations and the degree to which users trust the recommender, and also test the effect of giving the user the control on the fusing weights for explicit and implicit social network-based recommendations.

7. ACKNOWLEDGMENTS

This work has been supported by a PhD fellowship of the Institute of Public Administration, Riyadh, Saudi Arabia and by the NSERC Discovery Grant program.

8. References

[1] Avesani, P., Massa, P., and Tiella, R. A Trust-Enhanced Recommender System Application: Moleskiing. In *Proceedings of the 2005 ACM Symposium on Applied Computing*, (New York, USA, 2004), ACM Press, 1589–1593.

[2] Basu, C., Hirsh, H., Cohen, W.W., and Nevill-Manning, C. Technical Paper Recommendation: A Study in Combining Multiple Information Sources. *J Artif Int Res* 14 (2001), 231–252.

[3] Bellogín, A., Cantador, I., and Castells, P. A Study of Heterogeneity in Recommendations for a Social Music Service. In *Proceedings of the First International Workshop on Information Heterogeneity and Fusion in Recommender Systems*, (New York, NY, USA, 2010), ACM, 1–8.

[4] Bellogín, A., Cantador, I., Díez, F., Castells, & Chavarriaga, E. An empirical comparison of social, collaborative filtering, and hybrid recommenders. ACM Trans. Intell. Syst. Technol. 4 (2013), 14:1-14:29.

[5] Bostandjiev, S., O'Donovan, J., & Höllerer, T. TasteWeights: A Visual Interactive Hybrid Recommender System. In *Proceedings of the Sixth ACM Conference on Recommender Systems*, (New York, NY, USA, 2012) ACM, 35-42

[6] Burke, R. Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction*, 12,4 (2002), 331–370.

[7] Golbeck, J. FilmTrust: Movie Recommendations from Semantic Web-Based Social Networks. In *ISWC2005 Posters & Demonstrations*, (2005), PID–72.

[8] Groh, G., & Ehmig, C. Recommendations in Taste Related Domains: Collaborative Filtering vs. Social Filtering. In *Proceedings of the 2007 International ACM Conference on Supporting Group Work*, (New York, NY, USA, 2007), ACM, 127-136.

[9] Guy, I., Zwerdling, N., Carmel, D., Ronen, I., Uziel, E., Yogev, S., & Ofek-Koifman, S. Personalized Recommendation of Social Software Items Based on Social Relations. In *Proceedings of the Third ACM Conference on Recommender Systems*, (New York, NY, USA,2009), ACM, 53-60.

[10] Herlocker, J. L., Konstan, J. A., Terveen, L. G., John, & Riedl, T. (2004). Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.* 22, 1 (2004), 5-53

[11] Kautz, H., Selman, B., and Shah, M. Referral Web: Combining Social Networks and Collaborative Filtering. *Commun ACM* 40 (1997), 63–65.

[12] Ma, H., Yang, H., Lyu, M.R., and King, I. SoRec: Social Recommendation Using Probabilistic Matrix Factorization. In *Proceedings of the Seventeenth ACM Conference on Information and Knowledge Management*, (New York, NY, USA, 2008), ACM, 931–940.

[13] Massa, P., and Avesani, P. Trust-Aware Recommender Systems. In *Proceedings of the 2007 ACM Conference on Recommender Systems*, (New York, NY, USA, 2007), ACM, 17–24.

[14] McNee, S.M., Albert, I., Cosley, D., Gopalkrishnan, P., Lam, S.K., Rashid, A.M., Konstan, J.A., and Riedl, J. On the Recommending of Citations for Research Papers. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work*, (New York, NY, USA, 2002), ACM, 116–125.

[15] Mican, D., Mocean, L., & Tomai, N. (2012). Building a Social Recommender System by Harvesting Social Relationships and Trust Scores between Users. In W. Abramowicz, J. Domingue, & K. Węcel (Eds.), *Business Information Systems Workshops*, Springer Berlin Heidelberg. 1-12.

[16] Lee, D.H., and Brusilovsky, P. Improving Recommendations Using Watching Networks in a Social Tagging System. In *Proceedings of the 2011 iConference*, (New York, NY, USA, 2011), ACM, 33–39.

[17] Lee, D.H., and Brusilovsky, P. Interest Similarity of Group Members: The Case Study of CiteULike. In the *WebSci10: Extending the Frontiers of Society On-Line*, Raleigh, NC: USA, 2010).

[18] Lee, D. 2013. *Personalized Recommendations Based On Users' Information-Centered Social Networks*. Doctoral Dissertation. University of Pittsburgh.

[19] Liu, B. Informational Retrieval and Web Search. In *Web Data Mining: Exploring Hyperlinks, Contents and Usage Data*, (New York, NY, 2007), Springer, 183–236.

[20] Liu, F., and Lee, H.J. Use of Social Network Information to Enhance Collaborative Filtering Performance. *Expert Syst. Appl.* 37 (2010), 4772–4778.

[21] O'Donovan, J., and Smyth, B. Trust in Recommender Systems. In *Proceedings of the Tenth International Conference on Intelligent User Interfaces*, (New York, NY, USA, 2005), ACM, 167–174.

[22] Pera, M.S., and Ng, Y.-K. A Personalized Recommendation System on Scholarly Publications. In *Proceedings of the Twentieth ACM International Conference on Information and Knowledge Management*, (New York, NY, USA, 2011), ACM, 2133–2136.

[23] Pera, M.S., and Ng, Y.-K. Exploiting the Wisdom of Social Connections to Make Personalized Recommendations on Scholarly Articles. *J. Intell. Inf. Syst.* 42 (2014), 371–391.

[24] Tang, J., Hu, X., and Liu, H. Social Recommendation: *A Review. Soc. Netw. Anal. Min.* 3 (2013), 1113–1133.

[25] Torres, R., McNee, S.M., Abel, M., Konstan, J.A., and Riedl, J. Enhancing Digital Libraries with TechLens+. In *Proceedings of the Fourth ACM/IEEE-CS Joint Conference on Digital Libraries*, (New York, NY, USA, 2004), ACM, 228–236.

[26] Yuan, Q., Zhao, S., Chen, L., Liu, Y., Ding, S., Zhang, X., & Zheng, W. Augmenting collaborative recommender by fusing explicit social relationships. In *Workshop on Recommender Systems and the Social Web, with Recsys 2009*.