

Exploring Social Relations for Personalized Tag Recommendation in Social Tagging Systems

Kaipeng LIU[†], Binxing FANG^{††}, *Nonmembers*, and Weizhe ZHANG^{†a)}, *Member*

SUMMARY With the emergence of Web 2.0, social tagging systems become highly popular in recent years and thus form the so-called *folksonomies*. Personalized tag recommendation in social tagging systems is to provide a user with a ranked list of tags for a specific resource that best serves the user's needs. Many existing tag recommendation approaches assume that users are independent and identically distributed. This assumption ignores the social relations between users, which are increasingly popular nowadays. In this paper, we investigate the role of social relations in the task of tag recommendation and propose a personalized collaborative filtering algorithm. In addition to the social annotations made by collaborative users, we inject the social relations between users and the content similarities between resources into a graph representation of folksonomies. To fully explore the structure of this graph, instead of computing similarities between objects using feature vectors, we exploit the method of random-walk computation of similarities, which furthermore enable us to model a user's tag preferences with the similarities between the user and all the tags. We combine both the collaborative information and the tag preferences to recommend personalized tags to users. We conduct experiments on a dataset collected from a real-world system. The results of comparative experiments show that the proposed algorithm outperforms state-of-the-art tag recommendation algorithms in terms of prediction quality measured by precision, recall and NDCG.

key words: social tagging systems, personalized tag recommendation, social relations, collaborative filtering

1. Introduction

With the rapid development of Web 2.0, social tagging systems such as Delicious (<http://delicious.com/>) for sharing bookmarks, Flickr (<http://www.flickr.com/>) for sharing photos and CiteULike (<http://www.citeulike.org/>) for sharing academic publications, become highly popular in recent years. These systems allow collaborative users to submit shared resources and to annotate them with descriptive tags [1], forming the so-called *folksonomies*. We define a folksonomy as a structure $\mathbb{F} = (U, R, T, Y)$ consisting of i) a set U of *users*, ii) a set R of *resources*, iii) a set T of *tags*, and iv) the ternary relation between them, i.e. $Y \subseteq U \times R \times T$, called *annotations*. Though the resources being tagged can be webpages, images or videos, etc., we will focus on webpages in this paper unless otherwise mentioned.

A major reason for social tagging systems' immedi-

ate success is the highly easy-to-use user interface. Among other things, *tag recommendation* is by far the most important feature of the user interface. Tag recommendation refers to the automatic process of suggesting relevant and informative tags to an active user according to a specific resource based on historical information. The benefits of tag recommendation are twofold: enhancing user experience and enriching resource semantics. Though very important, the former is not emphasized by many previous studies. In fact, different people may have very different preferences while choosing tags for the same set of resources. Suggesting an unfamiliar tag to a user can be very frustrated even though it is used by many other users. Therefore, suggesting personalized tags to users that can properly serve their needs is of substantial importance to a sophisticated tag recommender.

Many existing tag recommendation approaches assume that users are independent and identically distributed (*i.i.d.*). Obviously, this assumption ignores the *social relations* between users, which are increasingly popular nowadays due to the common tendency for social tagging systems to encourage users to establish bonds of friendship between them. In Delicious, for example, a user can add another user sharing similar interests with him to his "network", by which automatically makes the former user become one of the latter user's "fans". By establishing bonds of friendship with others, one can build a *social network* to propagation his *social influence* for fulfilling his *social needs*.

The social relations between users can be very useful for the task of personalized tag recommendation. In a collaborative environment, it is more likely for users sharing similar interests to become friends. Friends are in turn heavily influenced by each other to form similar tagging behavior. This reinforcement relation inspires us that the tagging information of a user's friends can be leveraged to provide personalized tag recommendation for him. In other words, we are essentially taking *social recommendation* into account by exploring social relation. In fact, in real life, social recommendation plays an important role when people make choices. For example, when you ask a friend for a recommendation of a book to read or a movie to see, you are essentially soliciting a social recommendation. Previous studies have shown that, given both recommendations provided by friends and recommender systems, people tend to prefer those from the former, in terms of both quality and usefulness [2], [3].

In this paper, in addition to the social annotations made by collaborative users, we inject the social relations between

Manuscript received June 3, 2010.

Manuscript revised August 18, 2010.

[†]The authors are with the School of Computer Science and Technology, Harbin Institute of Technology, Harbin 150001, China.

^{††}The author is with the Key Laboratory of Network Science and Technology, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China.

a) E-mail: wzzhang@hit.edu.cn

DOI: 10.1587/transinf.E94.D.542

users into the relatedness data to leverage social recommendations from users' friends. We also symmetrically inject the content similarities between resources based on the intuition that similar resources tend to be annotated with similar tags. Specifically, we make the following contributions in this paper:

- We propose a graph representation of folksonomies that models the relatedness data including the social annotations, the social relations between users and the content similarities between resources.
- We exploit the method of random-walk computation of similarities to fully explore the structure of the graph representation.
- We develop a personalized collaborative filtering algorithm that combines both the collaborative information and the personalized tag preference.
- We conduct experiments on a dataset collected from a real-world system. The experimental results demonstrate the effectiveness of the proposed approach.

The remainder of this paper is organized as follows. Section 2 reviews previous studies on tag recommendation in social tagging systems; Sect. 3 describes the proposed graph representation of folksonomies; Sect. 4 introduces our approach for personalized social tag recommendation; Sect. 5 presents the experimental results on a real-world dataset; Sect. 6 concludes this work and presents some future works.

2. Related Work

Tag recommendation has attracted considerable attention in recent years. In this section, we review several major approaches for tag recommender systems. Based on the underlying data representation leveraged by different approaches, we generally divide tag recommenders into three categories, namely the content-based approaches, the collaborative filtering approaches and the link analysis-based approaches.

2.1 Content-Based Approaches

Content-based tag recommenders select tags that best describe the active resource. On one hand, the content of resources can be directly modeled to rank the relevant tags. The P-TAG algorithm [4] generated tags that were relevant to the textual content of the active resource as well as the documents residing on the user's desktop. Ralf Krestel et al. [5] used Latent Dirichlet Allocation (LDA) [6] to model the content of resources and used latent topics derived from social annotations to extend candidate tag set or recommend new tags to users. Ning Zhang et al. [7] employed a similar approach to model resources with the Author-Conference-Topic [8] model. Besides those text-based resources (e.g. webpages), the content of the other types of resources can be also used. For example, Lei Wu et al. [9] used the Visual Language Model (VLM) [10] to model the content of

tags in visual domain. They formulated the tag recommendation as a learning problem and used the features extracted from resources to learn an optimal combination of them. On the other hand, the content can be used to establish relations between resources or match resources to an external knowledge base, which are in turn integrated with the tag ranking algorithm. Yu-Ta Lu et al. [11] proposed a content-based method to enhance tag recommendation. Their approach was based on the simple assumption that similar resources tended to be annotated with similar tags. By propagating the probabilities for tags to annotate resources, tags were finally ranked with these probabilities. Ziyu Guan et al. [12] established relations among multi-type interrelated objects, i.e. users, resources and tags, via content similarity and proposed a graph-based ranking algorithm. Yang Song et al. [13] modeled the document distribution with a two-way Poisson Mixture Model (PMM), and classified the active resource with this model based on its posterior probabilities so that tags could be ranked with a within-cluster ranking function.

2.2 Collaborative Filtering Approaches

Among other things, *collaborative filtering* (CF) is by far the most popular technique employed by recommenders. Classical CF methods recommend items to users based on the preferences of similar users by exploiting the two-way relation of users and items [14]. Due to the ternary relational nature of social annotations, CF cannot be directly adapted. Previous studies [15], [16] employed a two-step method. Specifically, in the *collaborative* step, users sharing similar tag preferences with the active user were selected based on the between-user similarities. Then, in the *filtering* step, the best tags were selected based on the tagging information of these users. Robert Jäschke et al. [15] used two alternative 2-dimensional projections preserving user information for adapting CF to folksonomies by considering either resources or tags as objects. Marinho and Schmidt-Thieme [16] employed a similar approach and restricted the collaborative users to those who had annotated the active resource. Such methods cannot deal with the problem of *data sparsity* very well. For a new user who has posted only a few annotations, the similarities with existing users are hard to be properly calculated.

2.3 Link Analysis-Based Approaches

Another popular kind of tag recommenders are based on link analysis. Since the ternary relation among users, resources and tags can be expressed as a graph, link analysis-based approaches can be applied to such graph-based data representation. FolkRank [17] exploited the conceptual structures created by collaborative users. Inspired by the seminal PageRank [18] algorithm, FolkRank employed the underlying principle that a resource which was tagged with important tags by important users become important itself. The same held symmetrically for tags and users. This princi-

ple was modeled with a tripartite graph of vertices which mutually reinforced each other by spreading their weights. There were also other methods trying to exploit all the information of the ternary relations. Panagiotis Symeonidis et al. [19] used the Higher Order Singular Value Decomposition (HOSVD) technique [20] to directly perform latent semantic analysis and dimensionality reduction on the 3-dimensional social annotation data. Steffen Rendle et al. [21] proposed a method based on tensor factorization (TF). To directly optimize the tag ranking, they employed an optimization criterion and a learning algorithm, namely ranking with tensor factorization (RTF), for TF models.

3. Graph Representation of Folksonomies

The most intuitive way of expressing the ternary relations among users, resources and tags is to denote each object as a node in a graph, and the relations among them as edges. As shown in Fig. 1 (a), the nodes marked with $\{u_1, \dots, u_m\}$, $\{t_1, \dots, t_k\}$ and $\{r_1, \dots, r_n\}$ correspond to the users, tags and resources within a folksonomy, respectively. The edges can be weighted with a binary or a more sophisticated scheme. This graph representation was used by most previous studies. For example, for the task of tag recommendation, Robert Jäschke et al. [15] and Marinho et al. [16] adopted this graph representation to com-

pute the similarities between objects. They projected the ternary relation represented by this graph representation to a 2-dimensional space with a binary weighting scheme. For the task of ontology induction from folksonomies, Heymann and Garcia-Molina [22] employed a more sophisticated weighting scheme. To compute the similarities between tags, they used the number of times a tag was used to annotate a resource to weight the edge connecting them. More applications of this graph representation can be also found in [19], [23]–[25], etc.

In this paper, we extend this graph representation by augmenting the ternary relations with the social relations between users and the content similarities between resources. As shown in Fig. 1 (b), we add the edges between users to represent the social relations and the edges between resources to represent the content similarities. These connections are often ignored by previous studies. If properly used, however, they are very useful for the task of personalized tag recommendation. As discussed in Sect. 1, the social relations and content similarities can be leveraged to better model the user interest and resource content. By adding these edges to the graph representation of folksonomies, and employing a proper method to explore them, we can effectively improve the performance of tag recommenders. We will demonstrate the effectiveness of this graph representation via experimental evaluation in Sect. 5. Due to its ability to encode more relatedness data, we believe this graph representation can be potentially applied to more other tasks, such as resource (item) recommendation, user modeling, and user expertise ranking, etc.

As we can see in Fig. 1 (c), the $n \times n$ adjacency matrix of this graph representation, where $n = |U| + |R| + |T|$, can be divided into nine submatrices. In the following subsections, we describe how we fill these submatrices with proper values.

3.1 Modeling Social Annotations

The relations among users, resources and tags are established while users post resources to the system, and describe them with tags. There are two strategies which are often employed by existing studies to encode such relations. One alternative is to associate each pair of nodes (i, j) with a binary variable, i.e. $A_{i,j} = 1$ if i has co-occurred with j and 0 otherwise. The other alternative is to use the count of co-occurrences instead of the binary variable. Other than these strategies, we use another one based on the following convention: the more important the relation between two nodes, the larger the weight of the edge connecting them. Specifically, we use Jaccard's coefficient to measure the importance of relations, i.e.

$$A_{u,r}^{UR} = A_{r,u}^{RU} = |T(u) \cap T(r)| / |T(u) \cup T(r)|, \quad (1)$$

$$A_{u,t}^{UT} = A_{t,u}^{TU} = |R(u) \cap R(t)| / |R(u) \cup R(t)|, \quad (2)$$

$$A_{r,t}^{RT} = A_{t,r}^{TR} = |U(r) \cap U(t)| / |U(r) \cup U(t)|, \quad (3)$$

where $T(u) = \{t \mid \exists r: (u, r, t) \in Y\}$ and other sets are defined

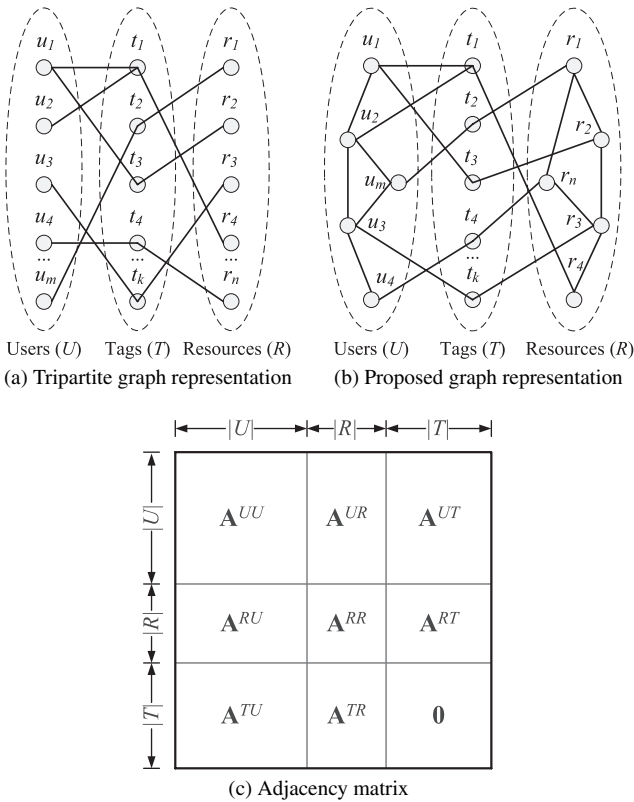


Fig. 1 Illustration of different graph representations of folksonomies: (a) the tripartite graph representation used by most existing studies, (b) the proposed graph representation used in this study, and (c) the adjacency matrix of the proposed graph representation.

in an analogous way. Since Jaccard's coefficient is indeed a similarity measure, we are essentially measuring the importance of a relation with the similarity between the corresponding nodes.

3.2 Modeling Social Relations

As mentioned in Sect. 1, a user's social network will affect his personal behavior. Based on this intuition, we inject the social relations between users into the graph representation to leverage the tagging information for users' friends. Most existing social tagging systems allow (encourage) users to establish bonds of friendships with others. These social relations are made public accessible if the permission from the user can be obtained. In Delicious, for example, we can acquire a user's social network by extracting users listed in the "network" and "fans" table of his profile. Notice that, in this specific case, these relations are not symmetric. The current user can only add another user into his "network", by which makes the former user become one of the latter user's "fans".

For a pair of users (u_i, u_j) , we define the weight of the edge connecting them as follows,

$$\mathbf{A}_{u_i, u_j}^{UU} = \text{trust}(u_i, u_j), \quad (4)$$

where $\text{trust}(u_i, u_j)$ is the trust level function specifying how much u_i trusts u_j . In real-world systems, the trust level function can be specified by users. For our Delicious dataset, we assume that these relations are symmetric and the trust levels are uniform distributed, i.e.

$$\text{trust}(u_i, u_j) = \begin{cases} 1, & \text{if } u_i \text{ is in } u_j\text{'s "network" or "fans" list,} \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

3.3 Modeling Content Similarities

It is quite intuitive that similar resources tend be annotated with similar tags. Despite the different personal tagging preferences, users usually annotate resources with tags that best describe the content of resources. Such tags remind users on what the resources really are and can help them to easily navigate to what they want to pick up.

For a pair of resources (r_i, r_j) , we use the content similarity between them as the weight for the edge connecting them, i.e.

$$\mathbf{A}_{r_i, r_j}^{RR} = \text{sim}(r_i, r_j), \quad (6)$$

where $\text{sim}(r_i, r_j)$ is a function that defines the similarity between r_i and r_j . Since computing the similarities between text-based documents are well studied in the literature, we are not going to discuss this issue in detail. Specifically, in our study, we use the Vector Space Model (VSM) [26] to model each resource with a TF-IDF based feature vector, and use cosine similarity to measure the content similarities between resources,

$$\text{sim}(r_i, r_j) = \cos(\mathbf{v}_i, \mathbf{v}_j) = (\mathbf{v}_i \cdot \mathbf{v}_j) / (\|\mathbf{v}_i\| \|\mathbf{v}_j\|), \quad (7)$$

where \mathbf{v}_i is the feature vector of r_i . To discard the negative impact of irrelevant resources, we eliminate the relations with a weight less than a predefined threshold θ , i.e. we assign a zero weight to the relation. For the experiment, we use an empirical value of $\theta = 0.01$.

4. Tag Recommendation

In this section, we employ the collaborative filtering technique for the task of tag recommendation. The basic idea of collaborative filtering is to predict tags for a given active user and resource on basis of the collaborative information available from similar users and resources. Thus, the notion of "similarity" plays an important role in such methods. After injecting all the relatedness data into the graph representation, we are now ready to explore the graph structure to compute the similarities between nodes.

4.1 Random-Walk Computation of Similarities

Most previous works calculated similarities between objects with their corresponding feature vectors, such as the row vectors extracted from the projection matrix, which was generated by projecting the ternary relation onto a 2-dimensional space [15], [16], [27]. Though simple and effective, it is not feasible for such methods to incorporate with additional information such as the social relations and the content similarities. Consider the social relations between users, for example, it is quite intuitive to say that two users are more similar if there are more direct or indirect relations between them. If we employ the methodology of representing a user's social relations as a feature vector, however, the similarity between two users only depends on the friends they have in common. Obviously, this method oversimplifies the structure of the user social networks—only the direct intermediate friends connecting different users are considered. Moreover, it ignores the direct relations between users, which are potentially more important than other relations. The similar analysis can be applied to other interrelations between other types of objects in folksonomies.

In this paper, we exploit the methodology of random-walk based computation of similarities between nodes. Similarity measures computed based on such models have the nice property of increasing while the number of paths connecting two nodes increases and when the "length" of any path decreases (i.e. when communication is facilitated) [28]. This property is quite appropriate for our graph representation. There are several existing methods for computing similarities (or dissimilarities instead) by exploiting the graph structure of relatedness data. Specifically, for two objects (i, j) in folksonomies, we consider the following similarity measures:

Average first-passage time (FPT) FPT [28], denoted by $\text{sim}(i, j) = m(i|j)$, is defined as the average number of steps for a random walker starting in state j to enter

state i . The FPT can be computed with:

$$\begin{cases} m(i|i) = 0 \\ m(i|j) = 1 + \sum_{k=1}^n \mathbf{P}_{j,k} m(i|k), \text{ for } j \neq i, \end{cases} \quad (8)$$

where $\mathbf{P}_{j,k} = \mathbf{A}_{j,k} / \sum_k \mathbf{A}_{j,k}$.

Average commute time (CT) CT, denoted by $\text{sim}(i, j) = n(i, j)$, is defined as the average number of steps for a random walker starting in state $i \neq j$ to enter state j for the first time and go back to i . That is,

$$n(i, j) = m(j|i) + m(i|j). \quad (9)$$

Pseudoinverse of the Laplacian matrix (\mathbf{L}^+) The symmetric Laplacian matrix \mathbf{L} of the graph is defined as $\mathbf{L} = \mathbf{D} - \mathbf{A}$, where $\mathbf{D}_{i,i} = \sum_j \mathbf{A}_{i,j}$. The Moore-Penrose pseudoinverse of \mathbf{L} , denoting by \mathbf{L}^+ , can be used as a similarity measure [28], i.e. $\text{sim}(i, j) = \mathbf{L}_{i,j}^+$. Moreover, let \mathbf{e} be a column vector made of 1s, \mathbf{L}^+ can be computed with the following formula (see [29], Ch. 10):

$$\mathbf{L}^+ = (\mathbf{L} - \mathbf{e}\mathbf{e}^T/n)^{-1} + \mathbf{e}\mathbf{e}^T/n. \quad (10)$$

Katz (Katz) Katz proposed a method of computing similarity matrix \mathbf{K} taking into account not only the number of direct links between nodes, but also the number of indirect links [30]. We use Katz's method to compute similarities between nodes, i.e. $\text{sim}(i, j) = \mathbf{K}_{i,j}$. The similarity matrix can be computed as

$$\mathbf{K} = \sum_{i=1}^{\infty} \alpha^i \mathbf{A}^i = (\mathbf{I} - \alpha \mathbf{A})^{-1} - \mathbf{I}, \quad (11)$$

where α is a positive constant s.t. $\alpha < \rho(\mathbf{A})^{-1}$, where $\rho(\mathbf{A})$ is the spectral radius of \mathbf{A} . For the experiment, we systematically varied the value of $\alpha = (0.05, 0.10, \dots, 0.95) \cdot \rho(\mathbf{A})^{-1}$ to find the best setting, namely $\alpha = 0.05 \cdot \rho(\mathbf{A})^{-1}$.

Matrix-forest-based algorithm (MFA) Chebotarev and Shamis proposed in [31], [32] a similarity matrix \mathbf{T} that has an interest interpretation in terms of the matrix forest theorem (refer to [31], [32] for details). It can be shown that this matrix is a similarity measure, i.e. $\text{sim}(i, j) = \mathbf{T}_{i,j}$, having the natural properties such as triangular property. The similarity matrix is

$$\mathbf{T} = (\mathbf{I} + \mathbf{L})^{-1}. \quad (12)$$

Note that some of the above measures, such as FPT and CT, are indeed dissimilarity measures. While applying such measures, the ranking list needs to be reversed to get the final results.

4.2 Collaborative Filtering

Two variants of collaborative filtering (CF) method, namely user-based CF (User-CF) and resource-based CF (Resource-CF), are well studied in the literature [16], [27]. User-CF is aiming at finding the similar users with the active user

and recommends tags based on the tagging information of these similar users. Specifically, for a given active user u and resource r , the top- K similar users who have tagged r are first selected to form the K -neighborhood of u ,

$$N_U^K(u, r) = \arg \max_{u' \in U(r)}^K \text{sim}(u, u'), \quad (13)$$

where $U(r) = \{u | \exists t : (u, r, t) \in Y\}$, and then for a given number of suggested tags N , the final suggested tags $\tilde{T}^N(u, r)$ is generated as follows,

$$\tilde{T}^N(u, r) = \arg \max_{t \in T}^N \sum_{u' \in N_U^K(u, r)} \text{sim}(u, u') \delta(u', r, t), \quad (14)$$

where $\delta(u, r, t) = 1$ if $(u, r, t) \in Y$ and 0 else. Resource-CF is analogous to User-CF. Specifically, we first select the K -neighborhood of r ,

$$N_R^K(u, r) = \arg \max_{r' \in R(u)}^K \text{sim}(r, r'), \quad (15)$$

where $R(u) = \{r | \exists t : (u, r, t) \in Y\}$, and suggest tags with

$$\tilde{T}^N(u, r) = \arg \max_{t \in T}^N \sum_{r' \in N_R^K(u, r)} \text{sim}(r, r') \delta(u, r', t). \quad (16)$$

4.3 Personalized Tag Recommendation

In addition to the CF methods described above, we propose a personalized CF method (Personalized-CF) to incorporate each user's tag preference with the collaborative information from other users. In fact, a CF-based tag recommender may fail to predict appropriate tags for a given user for two reasons. The first reason is that, the active user's tag preference is different from the others. For example, while many users tag a webpage about mobile phone with `mobile-phone`, a few others may prefer to tag it with `cell-phone`. The second reason is that, the active user's tag vocabulary is different from the others. For example, a user likes to tag webpages about books with tags such as `toread` and `tobuy`, which are hard to predict from the collaborative information from the others. We propose to solve these problems by incorporating the collaborative information with the personalized tag preferences.

To obtain the personalized tag preference for a user, we should know how much he "favors" a tag. By exploiting the methodology of random-walk computation of similarities, we are able to compute similarities between heterogeneous objects. Therefore, each user's tag preference can be represented by the similarities between this user and all the tags. Personalized-CF generates the recommended tags as follows,

$$\tilde{T}^N(u, r) = \arg \max_{t \in T}^N (1 - \lambda) \cdot c(u, r, t) + \lambda \cdot p(u, t), \quad (17)$$

where $\lambda \in [0, 1]$ is a predefined constant, $c(u, r, t)$ is the collaborative information that is personalized with the user's

tag preference,

$$c(u, r, t) = p(u, t) \cdot \sum_{u' \in N_{ij}^K(u, r)} \text{sim}(u, u') \delta(u', r, t), \quad (18)$$

and $p(u, t)$ is the user's tag preference,

$$p(u, t) = \text{sim}(u, t). \quad (19)$$

With Eq. (17), we aim to solve the problem of different tag preferences with Eq. (18) and different tag vocabulary with Eq. (19). For the experiments, we systematically varied the value of $\lambda = (0.05, 0.10, \dots, 0.95)$ to find the best setting, namely $\lambda = 0.15$.

5. Experiments

We conduct experiments on a dataset collected from a real-world system to evaluate the performance of the proposed approach. In the following subsections, we first introduce the dataset used in the experiments, then describe the evaluation methodology and finally present the experimental results.

5.1 Dataset

We conduct experiments on the Delicious dataset. Starting at Dec 2007, we have crawled thousands of web pages from Delicious and extracted post information including user, resource, post date and corresponding tags. The preprocess of the raw dataset was done in two steps. First, we extracted the social relations between users, which were listed in a user's "network" and "fans" table. Totally 313,151 social relations were extracted in this dataset. We eliminated those users with no social connection, along with those annotations made by them. One may argue that, eliminating such users in the raw dataset may result in a bias towards the proposed method. Since we are investigating the additional value to incorporate with such social relations between users for the task of tag recommendation, however, it does not make sense to perform the comparison on those users with no social relation at all. Second, we computed the p -core [33] at level k by filtering out the users, resources and tags which occurred less than k times. By p -core computation, those inactive users, unpopular resources and idiosyncratic tags were eliminated from the dataset. According to the result of preliminary experiment, the value of k had little impact on the relative performance of different algorithms. We used $k = 10$ for the main experiment. There were 8,133 users, 18,350 resources, 5,966 tags and 3,546,718 annotations in the preprocessed dataset.

5.2 Evaluation Methodology

We used the same evaluation protocol as previous studies [15], [21], [27], which was a variant of the leave-one-out hold-out estimation [34]. To generate one split of the

dataset, for each user in the dataset, we removed all the annotations associated with a random-selected resource. All the removed annotations formed the test set Y_{test} and the remaining annotations formed the training set $Y_{\text{train}} = Y \setminus Y_{\text{test}}$. Then we used the training set as input to tag recommenders and used the test set to evaluate performance of them.

All the experiments were performed in a two-stage manner. In the *preliminary experiment*, we run all the algorithms once with one split of the dataset to tune the parameters of all the algorithms (the results are not shown in this paper). In the *main experiment*, all the algorithms were repeated for 10 times with 10 different splits of the dataset and the average results were reported.

To evaluate the prediction quality of a tag recommender, we use the standard metrics of precision, recall and NDCG. For a given number N of suggested tags and a pair of active user u and resource r , these metrics are defined as

$$\begin{aligned} \text{Precision} &= \frac{|\tilde{T}^N(u, r) \cap T(u, r)|}{N}, \\ \text{Recall} &= \frac{|\tilde{T}^N(u, r) \cap T(u, r)|}{|T(u, r)|}, \\ \text{NDCG} &= \frac{1}{I_Q} \sum_{j=1}^N \frac{2^{\text{rel}(u, r, j)} - 1}{\log(1 + j)}, \end{aligned}$$

where I_Q is a normalization constant to make a perfect ranking result achieve a NDCG value of 1. We use a binary relevance score, i.e. $\text{rel}(u, r, j) = 1$ if the j th tag in $\tilde{T}^N(u, r)$ is also in $T(u, r)$ and 0 otherwise. For each metric, the results over all the user-resource pairs were averaged to get the final results.

5.3 Comparison of Similarity Measures

We begin by comparing the five similarity measures listed in Sect. 4.1. We set $K = 100$ and $N = 10$ for the three CF methods, namely User-CF, Resource-CF and Personalized-CF. Notice that the setting of these parameters has little impact on the relative performance of all the algorithms. The results are shown in Fig. 2 and Fig. 3.

From Fig. 2 and Fig. 3 we can see that MFA achieves the best performance for User-CF and Personalized-CF, while L^+ for Resource-CF. L^+ seems to lack stability, however, since the results are very sensitive to the method used. The dissimilarity measures, i.e. FPT and CT, are clearly less efficient than other measures.

We also compare the performance of all the three CF methods with the best similarity measures in Fig. 2 (d) and Fig. 3 (d). It is shown that Personalized-CF performs the best while Resource-CF performs the worst. The fact that the two user-based methods outperform the resource-based method indicates that the collaborative information from similar users is more efficient than that from similar resources. Comparing Personalized-CF with User-CF, it is clearly shown that the former can recommend more appropriate tags by taking users' tag preferences into account.

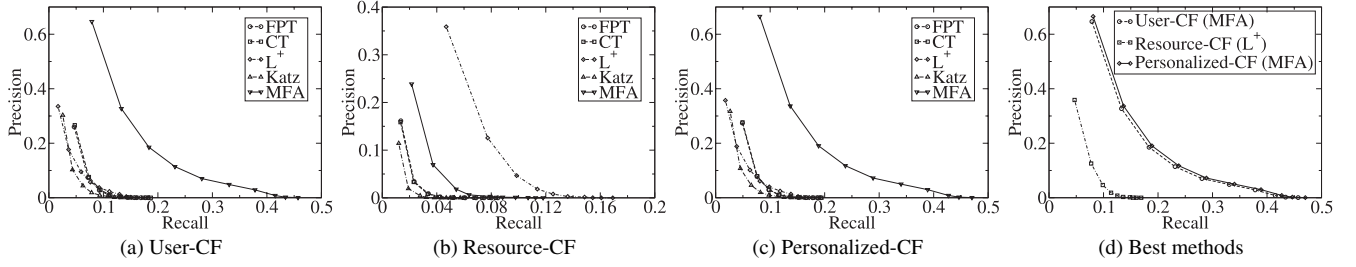


Fig. 2 Comparison of different similarity in Recall/Precision for (a) User-CF, (b) Resource-CF, and (c) Personalized-CF. We also compare in (d) the methods achieving the best performance.

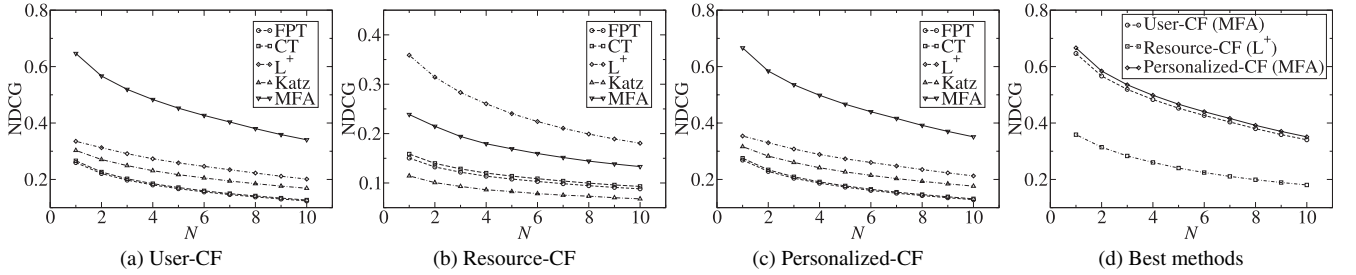


Fig. 3 Comparison of different similarity measures in NDCG for (a) User-CF, (b) Resource-CF, and (c) Personalized-CF. We also compare in (d) the methods achieving the best performance.

5.4 Comparison of Recommendation Algorithms

We compare the best algorithm, namely the Personalized-CF with MFA, with other state-of-the-art tag recommendation algorithms including collaborative filtering methods employing feature-based computation of similarity [16], [27] and mutual information (MI) computation of similarity [24], and FolkRank [17].

The collaborative filtering (CF) methods are analogous to the User-CF and Resource-CF algorithm described in Sect. 4.2, except for the similarity computation method. For user-based CF, the feature-based similarity is computed with cosine similarity based on the projection aggregation method [24]. This similarity computation method was used in [16], [27]. The ternary relation Y is reduced to a lower dimensional space with two 2-dimensional projections: $\pi^{UT}(Y) \in 0, 1^{|U| \times |T|}$ with $\pi^{UT}(Y)_{u,t} = 1$ if there exists $r \in R$ s.t. $(u, r, t) \in Y$ and 0 otherwise, and $\pi^{UR}(Y) \in 0, 1^{|U| \times |R|}$ with $\pi^{UR}(Y)_{u,r} = 1$ if there exists $t \in T$ s.t. $(u, r, t) \in Y$ and 0 otherwise. The similarities between users are computed with $\text{sim}(u, u') = \text{sim}(\pi(Y)_u, \pi(Y)_{u'}) = (\pi(Y)_u \cdot \pi(Y)_{u'}) / (\|\pi(Y)_u\| \cdot \|\pi(Y)_{u'}\|)$. The mutual information-based similarity is computed with mutual information based on the distributional aggregation method [24]. The ternary relation Y is reduced to a 2-dimensional space with two distributional weighting scheme: $\tau^{UT}(Y) \in \mathbb{N}^{|U| \times |T|}$ with $\tau^{UT}(Y)_{u,t} = |R(u, t)|$ and $\tau^{UR}(Y) \in \mathbb{N}^{|U| \times |R|}$ with $\tau^{UR}(Y)_{u,r} = |T(u, r)|$, where $R(u, t)$ is the set of resources associated with user u and tag t , and $T(u, r)$ is set of tags associated with user u and resource r . The similarities between users are computed with $\text{sim}(u, u') =$

$\sum_{x \in \tau(Y)_u} \sum_{x' \in \tau(Y)_{u'}} p(x, x') \log(p(x, x') / (p(x)p(x')))$, where $x \in \tau(Y)_u$ means that $\tau(Y)_{u,x} \neq 0$. The probabilities are computed with $p(x) = \sum_u \tau(Y)_{u,x} / \sum_{u,y} \tau(Y)_{u,y}$ and $p(x, x') = \sum_u \min(\tau(Y)_{u,x}, \tau(Y)_{u,x'}) / \sum_{u,y} \tau(Y)_{u,y}$. For resource-based CF, the similarities between resources are computed in an analogous way.

The FolkRank algorithm first computes the adapted PageRank with $\mathbf{w} = d\mathbf{A}\mathbf{w} + (1-d)\mathbf{p}$, where \mathbf{A} is the adjacency matrix (refer to [17] for detail) and \mathbf{p} is the preference vector. Let \mathbf{w}^0 and \mathbf{w}^1 be the fix points from the iteration with $d = 1$ and $d < 1$. For the experiment, we use the same setting of $d = 0.7$ as the origin work [15]. The final score vector is computed with a differential method, i.e. $\mathbf{w} = \mathbf{w}^1 - \mathbf{w}^0$. While applying FolkRank to the task of tag recommendation, for the preference vector \mathbf{p} , we assign a higher weight of $1 + |U|$ and $1 + |R|$ to the active user and resource, while a weight of 1 to the other objects. Tags with the highest FolkRank score are recommended.

We compare the CF method based on graph representation with and without additional information such as social relations and content similarities in Fig. 4 (a) and Fig. 4 (b). As we can see from these figures, algorithms augmented with additional information outperform the others for both user-based and resource-based CF methods. This clearly shows that injecting additional information and performing random-walk computation of similarities effectively improves the performance of tag recommenders..

The comparison between User-CF and Resource-CF conforms to the result in Sect. 5.3, despite we are now using different similarity measures. This confirms us that collaborative information from similar users is more efficient than that of similar resources. For User-CF, the π^{UT} -based

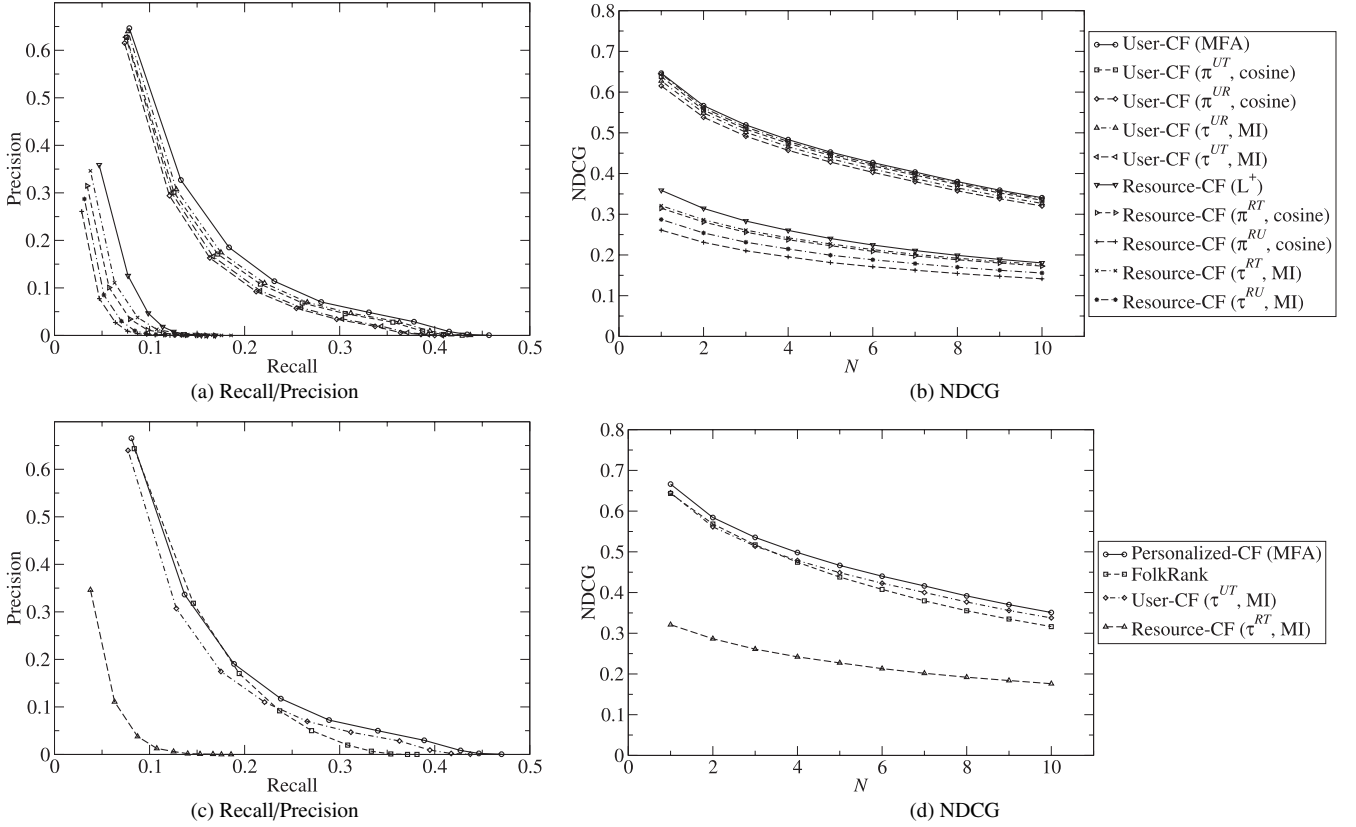


Fig. 4 Comparison of different tag recommenders. In (a) and (b), the CF methods based on graph representation w/ and w/o additional information. In (c) and (d), the proposed Personalized-CF is compared with other state-of-the-art methods such as collaborative filtering based methods and FolkRank.

cosine similarity and τ^{UT} -based MI similarity outperforms the π^{UR} -based cosine similarity and τ^{UR} -based MI similarity. This indicates that the collaborative information from users sharing similar tag preference is more efficient than that from users sharing similar resource preference. Similar result can be also observed while comparing the two versions of Resource-CF. For the same aggregation object type, the MI similarity outperforms the cosine similarity slightly.

We also compare Personalized-CF with the CF methods based on MI similarity and the other methods such as FolkRank in Fig. 4(c) and Fig. 4(d). Comparing these cosine similarity based CF methods with Personalized-CF, we can see that Personalized-CF can achieve a better performance.

It is interesting to note that the performance of FolkRank is very good when the number of recommended tags (N) is small, but drops quickly while N increases. For FolkRank, importance can be only propagated through the ternary relations between users, resources and tags. For Personalized-CF, however, the social relations between users and the content similarities between resources can be also exploited, which results in a better performance.

6. Conclusions and Future Works

In this paper we presented our approach to augment person-

alized tag recommenders with the social relations between users and the content similarities between resources and developed a personalized collaborative filtering algorithm. We showed that, with proper similarity measures, we could effectively improve the performance of tag recommenders by incorporating with such additional information. The most important findings in the experimental analysis were:

- Among all the similarity measures, MFA gave the best performance for User-CF and Personalized-CF, while L^+ for Resource-CF.
- The proposed Personalized-CF algorithm with MFA as similarity measure achieved the best performance of all the algorithms compared.
- The collaborative information from similar users was more efficient than that from similar resources.

All these findings could provide us with useful information for building sophisticated tag recommenders.

The future works are twofold. On the one hand, we are examining more similar measures to find the most efficient one. On the other hand, we are trying to incorporate more information into tag recommenders, such as hyperlinks between webpages, semantic and lexical relations between tags, etc., to improve the performance of tag recommendation algorithm.

References


- [1] S.A. Golder and B.A. Huberman, "Usage patterns of collaborative tagging systems," *J. Inform. Sci.*, vol.32, no.2, pp.198–208, 2006.
- [2] I. Guy, N. Zwerdling, D. Carmel, I. Ronen, E. Uziel, S. Yogeve, and S. Ofek-Koifman, "Personalized recommendation of social software items based on social relations," *Proc. 3rd ACM Conf. Recommender Systems (RecSys'09)*, pp.53–60, New York, NY, USA, Oct. 2009.
- [3] P. Bonhard and A. Sasse, "“Knowing me, knowing you” - Using profiles and social networking to improve recommender systems," *BT Technol. J.*, vol.24, no.3, pp.84–98, 2006.
- [4] P.A. Chirita, S. Costache, W. Nejdl, and S. Handschuh, "P-tag: Large scale automatic generation of personalized annotation tags for the web," *Proc. 16th Int. Conf. World Wide Web (WWW'07)*, pp.845–854, Banff, Alberta, Canada, May 2007.
- [5] R. Krestel, P. Fankhauser, and W. Nejdl, "Latent dirichlet allocation for tag recommendation," *Proc. 3rd ACM Conf. Recommender Systems (RecSys'09)*, pp.61–68, New York, NY, USA, Oct. 2009.
- [6] D.M. Blei, A.Y. Ng, M.I. Jordan, and J. Lafferty, "Latent dirichlet allocation," *J. Mach. Learn. Res.*, vol.3, pp.993–1022, 2003.
- [7] N. Zhang, Y. Zhang, and J. Tang, "A tag recommendation system for folksonomy," *Proc. 2nd ACM Workshop on Social Web Search and Mining (SWSM'09)*, pp.9–16, Hong Kong, China, Nov. 2009.
- [8] J. Tang, R. Jin, and J. Zhang, "A topic modeling approach and its integration into the random walk framework for academic search," *Proc. 8th IEEE Int. Conf. Data Mining (ICDM'08)*, pp.1055–1060, Pisa, Italy, Dec. 2008.
- [9] L. Wu, L. Yang, N. Yu, and X.S. Hua, "Learning to tag," *Proc. 18th Int. Conf. World Wide Web (WWW'09)*, pp.361–370, Madrid, Spain, April 2009.
- [10] L. Wu, M. Li, Z. Li, W.Y. Ma, and N. Yu, "Visual language modeling for image classification," *Proc. Int. Workshop on Multimedia Information Retrieval (MIR'07)*, pp.115–124, Augsburg, Bavaria, Germany, Sept. 2007.
- [11] Y.T. Lu, S.I. Yu, T.C. Chang, and J.Y.J. Hsu, "A content-based method to enhance tag recommendation," *Proc. 21st Int. Joint Conf. Artificial Intelligence (IJCAI'09)*, pp.2064–2069, Pasadena, California, USA, July 2009.
- [12] Z. Guan, J. Bu, Q. Mei, C. Chen, and C. Wang, "Personalized tag recommendation using graph-based ranking on multi-type interrelated objects," *Proc. 32nd Int. ACM SIGIR Conf. Research and Development in Information Retrieval (SIGIR'09)*, pp.540–547, Boston, MA, USA, July 2009.
- [13] Y. Song, Z. Zhuang, H. Li, Q. Zhao, J. Li, W.C. Lee, and C.L. Giles, "Real-time automatic tag recommendation," *Proc. 31st Ann. Int. ACM SIGIR Conf. Research and Development in Information Retrieval (SIGIR'08)*, pp.515–522, Singapore, Singapore, July 2008.
- [14] B. Sarwar, G. Karypis, J. Konstan, and J. Reidl, "Item-based collaborative filtering recommendation algorithms," *Proc. 10th Int. Conf. World Wide Web (WWW'01)*, pp.285–295, Hong Kong, China, May 2001.
- [15] R. Jäschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme, "Tag recommendations in folksonomies," *Proc. 11th European Conf. Principles and Practice of Knowledge Discovery in Databases (PKDD'07)*, pp.506–514, Warsaw, Poland, Sept. 2007.
- [16] L.B. Marinho and L. Schmidt-Thieme, "Collaborative tag recommendations," *Proc. 31st Ann. Conf. German Classification Society (GfKI'07)*, pp.533–540, Freiburg, Germany, March 2007.
- [17] A. Hotho, R. Jäschke, C. Schmitz, and G. Stumme, "Information retrieval in folksonomies: Search and ranking," *Proc. 3rd European Semantic Web Conference (ESWC'06)*, pp.411–426, Budva, Montenegro, June 2006.
- [18] L. Page, S. Brin, R. Motwani, and T. Winograd, "The pagerank citation ranking: Bringing order to the web," *Technical Report 1999-66*, Stanford InfoLab, Nov. 1999.
- [19] P. Symeonidis, A. Nanopoulos, and Y. Manolopoulos, "Tag recommendations based on tensor dimensionality reduction," *Proc. 2008 ACM Conf. Recommender Systems (RecSys'08)*, pp.43–50, Lausanne, Switzerland, Oct. 2008.
- [20] L.D. Lathauwer, B.D. Moor, and J. Vandewalle, "A multilinear singular value decomposition," *SIAM J. Matrix Anal. Appl.*, vol.21, no.4, pp.1253–1278, 2000.
- [21] S. Rendle, L.B. Marinho, A. Nanopoulos, and L. Schmidt-Thieme, "Learning optimal ranking with tensor factorization for tag recommendation," *Proc. 15th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD'09)*, pp.727–736, Paris, France, June 2009.
- [22] P. Heymann and H. Garcia-Molina, "Collaborative creation of communal hierarchical taxonomies in social tagging systems," *Tech. Rep. 2006-10*, Stanford University, April 2006.
- [23] A. Shepitsen, J. Gemmell, B. Mobasher, and R. Burke, "Personalized recommendation in social tagging systems using hierarchical clustering," *Proc. 2008 ACM Conference on Recommender Systems (RecSys'08)*, pp.259–266, Lausanne, Switzerland, Oct. 2008.
- [24] B. Markines, C. Cattuto, F. Menczer, D. Benz, A. Hotho, and G. Stumme, "Evaluating similarity measures for emergent semantics of social tagging," *Proc. 18th International Conference on World Wide Web (WWW'09)*, pp.641–650, Madrid, Spain, April 2009.
- [25] H. Halpin, V. Robu, and H. Shepherd, "The complex dynamics of collaborative tagging," *Proc. 16th International Conference on World Wide Web (WWW'07)*, pp.211–220, May 2007.
- [26] G. Salton, A. Wong, and C.S. Yang, "A vector space model for automatic indexing," *Commun. ACM*, vol.18, no.11, pp.613–620, 1975.
- [27] R. Jäschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme, "Tag recommendations in social bookmarking systems," *AI Commun.*, vol.21, no.4, pp.231–247, 2008.
- [28] F. Fouss, A. Pirotte, J.M. Renders, and M. Saerens, "Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation," *IEEE Trans. Knowl. Data Eng.*, vol.19, no.3, pp.355–369, 2007.
- [29] C. Rao and S. Mitra, *Generalized inverse of a matrix and its applications*, John Wiley & Sons, New York, 1971.
- [30] L. Katz, "A new status index derived from sociometric analysis," *Psychometrika*, vol.18, no.1, pp.39–43, 1953.
- [31] P. Chebotarev and E. Shamis, "The matrix-forest theorem and measuring relations in small social groups," *Autom. Remote Control*, vol.58, no.9, pp.1505–1514, 1997.
- [32] P. Chebotarev and E. Shamis, "On proximity measures for graph vertices," *Autom. Remote Control*, vol.59, no.10, pp.1443–1459, 1998.
- [33] V. Batagelj and M. Zaversnik, "Generalized cores," *CoRR*, vol.cs.DS/0202039, 2002.
- [34] J.L. Herlocker, J.A. Konstan, L.G. Terveen, and J.T. Riedl, "Evaluating collaborative filtering recommender systems," *ACM Trans. Inf. Syst.*, vol.22, no.1, pp.5–53, 2004.



Kaipeng Liu is currently a Ph.D. candidate at Harbin Institute of Technology, China. He received his MS and BS degrees in computer science and technology from the Harbin Institute of Technology in 2006 and 2004, respectively. His research interests include information retrieval, data mining and machine learning.



Binxing Fang received his Ph.D. degree in computer architecture from Harbin Institute of Technology in 1989, his MS degree in computer science and technology from the Tsinghua University in 1984, and his BS degree in computer science and technology from Harbin Institute of Technology in 1981. He is currently a member of Chinese Academy of Engineering. His current research interests include information security, information retrieval, and distributed systems.



Weizhe Zhang is doctor, associate professor. Research fields include: Network computing, cluster computing, parallel and distributed system. Undertaking projects include: National Natural Science Foundation of China, 863 High-Tech Research Plan of China, Province and Ministry Science Foundation projects. Published over 30 papers in journals and international conferences.