Knowledge-Based Systems 105 (2016) 83–95

Contents lists available at ScienceDirect

Knowledge-Based Systems

journal homepage: www.elsevier.com/locate/knosys

A framework for diversifying recommendation lists by user interest expansion

Zhu Zhang[a], Xiaolong Zheng[a,*], Daniel Dajun Zeng[a,b]

[a] Institute of Automation, Chinese Academy of Sciences, Beijing, China, 100190
[b] Department of Management Information Systems, The University of Arizona, Tucson, USA

ARTICLE INFO

Article history:
Received 24 July 2015
Revised 3 May 2016
Accepted 7 May 2016
Available online 9 May 2016

Keywords:
Recommender systems
Collaborative filtering
Diversity
Interest expansion
Social tagging system

ABSTRACT

Recommender systems have been widely used to discover users' preferences and recommend interesting items to users during this age of information overload. Researchers in the field of recommender systems have realized that the quality of a top-N recommendation list involves not only relevance but also diversity. Most traditional recommendation algorithms are difficult to generate a diverse item list that can cover most of his/her interests for each user, since they mainly focus on predicting accurate items similar to the dominant interests of users. Additionally, they seldom exploit semantic information such as item tags and users' interest labels to improve recommendation diversity. In this paper, we propose a novel recommendation framework which mainly adopts an expansion strategy of user interests based on social tagging information. The framework enhances the diversity of users' preferences by expanding the sizes and categories of the original user-item interaction records, and then adopts traditional recommendation models to generate recommendation lists. Empirical evaluations on three real-world data sets show that our method can effectively improve the accuracy and diversity of item recommendation.

1. Introduction

With the fast growth of web resources such as web pages, music, movies, etc. (referred as items), lots of Internet users are confused by how to select their favorite items from huge amount of items. Since the mid-1990 s, Recommender Systems (RS), as an effective information filtering tool to discover users' preferences and recommend interesting items to users, have drawn much attention from academia and industry [1,2]. In real-world applications, RS have been used to recommend web pages, books, music, movies, TV programs, places of interest, conferences, courses, etc. Generally, RS applications can be categorized into eight main domains: e-government [3], e-business [4], e-commerce/e-shopping [5], e-library [6], e-learning [7], e-tourism [8], e-resource services [9], and e-group activities [10], and the application platforms include Web, TV, radio, and mobile devices [11]. Additionally, as users of RS applications are individuals or groups, the corresponding RS can be classified into individual-based or group-based.

No matter the input data of RS are explicit (e.g., rating and like/dislike) [12,13] or implicit (e.g., clicking and purchase) [14,15], most of RS usually provide a top-N item list for each user. A high-quality top-N recommendation list should have not only high relevance but also high diversity [16]. In this paper, diversity generally indicates the pairwise differences of the recommended items in a top-N list [17], and it has an important impact on users' experiences. For instance, if a user indicates he watched three movies directed by the famous director James Cameron and a few other movies on a movie review website, and then the RS of this site give this user a top-10 recommendation list in which nearly all of the movies are directed by this director. It is obvious that the list is not so appealing due to its monotony.

Traditional RS, including two main categories: latent factor models [14,18] and neighborhood-based models [12,19], are hard to generate a diverse item list that can cover most of one user's interests, since those algorithms mainly focus on predicting accurate items that are very similar to the dominant interests of users [20]. They are also difficult to estimate users' extensive interests due to the insufficiency of original item records' number and categories. In addition, most of traditional RS mainly focus on accuracy in terms of performance evaluation, and the involved measure metrics include precision and recall in information retrieval, and MAE and RMSE in rating prediction. To solve these problems, some studies have attempted to develop several metrics to quantify the diversity of recommendation results [17,20,21] and further to propose new recommendation models for improving the diversity. Among these studies, some of these models are composed

* Corresponding author.
E-mail addresses: zhu_zhang@foxmail.com (Z. Zhang), xiaolong.zheng@ia.ac.cn (X. Zheng), dajun.zeng@ia.ac.cn (D.D. Zeng).

http://dx.doi.org/10.1016/j.knosys.2016.05.010
0950-7051/

of a relevance-oriented part and a diversity-oriented part [21], while some other models take into account relevance and diversity in a unified framework [20,22]. The two-part hybrid modes can incorporate traditional recommendation algorithms and leverage their advantages, while the unified diversification frameworks are more compact and principled in terms of formulation. However, these models don't consider expanding users' interests based on item usage records to improve recommendation diversity, although the expansion of users' interests in the process of building recommendation models may be beneficial to diversity intuitively. Additionally, the recommendation models in these studies only exploit users' item records to build users' interest profiles, so they usually ignore some semantic information such as item tags and users' interest labels. However, tags have the potential to help improve the accuracy and diversity of recommendations, since tags can represent user interests and item features effectively. Some researchers have exploited semantic tag information and the associations among users, tags, and items of the existing Social Tagging Systems (STSs) such as Delicious[1], BibSonomy[2], and Flickr[3], to improve recommendation accuracy. In general, the low diversity of recommendation results is still one of the major research problems in current RS studies.

To tackle this problem, we propose a novel user-interest-expansion-based recommendation framework by exploiting social tagging information, which can be adopted to enhance the diversity of recommendation lists. The proposed interest expansion strategy includes three key steps. Firstly, we represent users' features with users' tag records and compute the similarity between tags by item-tag interaction information. Then we expand users' initial tag records based on tag similarity. Secondly, we select new items with proper popularity from the items relating to each user's expanded tag sets, and add those selected items into users' initial item sets. Finally, we treat the expanded users' item sets as the input of traditional item recommendation algorithms, and then we obtain personalized and diverse recommendation lists.

The main contributions of this paper can be summarized as follows:

1. The first innovation of this study is the proposed recommendation diversification framework which mainly utilizes user interest expansion based on social tagging information. To the best of our knowledge, few studies consider enhancing recommendation diversity by exploring social tagging information in the literature concerning RS and STS. The idea of expanding user interests for the improvement of recommendation diversity is also innovative. One advantage of this innovative framework is that it can leverage different kinds of advanced traditional RS within our framework to ensure high recommendation accuracy. Another advantage is that adopting social tagging information for user interest expansion can facilitate the joint improvement of recommendation accuracy and diversity.
2. The second innovation is that we develop a novel diversity metric named tag coverage (i.e., TagCov in the following tables and figures) in the experimental evaluation. The metric can evaluate the diversity of users' interests efficiently, which increase the diversity measure metrics regarding RS evaluation.
3. Experimental results on three real-world data sets demonstrate that our approach can outperform existing ones. Especially, our approach can improve the diversity and accuracy of item recommendations more efficiently than existing ones.

The rest of this paper is organized as follows. We present related work in Section 2 and introduce the proposed research method in Section 3. Then the experimental evaluation is presented in Section 4. In Section 5, we conclude this paper and present future work.

2. Related work

In general, the related work can be grouped into two parts: recommendation diversity and tag-based item recommendation. The former is related to how to measure and improve the diversity of a recommendation list for each user, and the latter is about how to enhance the accuracy of item recommendations by leveraging tag information.

2.1. Recommendation diversity

Accuracy is not the only measure dimension of recommendation results. The importance of diversity in terms of RS performance evaluation had drawn much attention in the early studies [16,23,24]. For example, McNee et al., [16] argued that accuracy is not always useful and proposed some new user-centric directions for recommendation evaluation. According to [25], the diversity measurement of recommendations can be classified into two types: individual diversity and aggregate diversity. Individual diversity emphasizes the recommendation diversity from an individual user's perspective, which can usually be calculated by an average dissimilarity of all pairs of items recommended to a given user [17,20,21,23]. However, aggregate diversity emphasizes the diversity of recommendations across all users, which can be measured by some metrics (e.g., coverage and diversity-in-top-N) [25]. In this paper, we only focus on the individual diversity of recommendations. For convenience, we will refer to individual diversity as diversity throughout the paper, unless explicitly specified otherwise.

Different diversity metrics were proposed in previous studies. Vargas et al., [17] summarized the state-of-the-art diversity metrics and presented a formalized framework which can generalize these metrics. The common computational method of diversity is provided in the above definition of individual diversity, and the computation of dissimilarity is often based on the semantic information of items or user-item interaction records. In addition, some strategies were proposed to improve the recommendation diversity. For instance, Ziegler et al., [21] proposed a novel model that re-ranks the recommendation list generated by traditional Collaborative Filtering (CF) according to the topic diversity metric, and the online experiments demonstrated that the recommendation accuracy decreased but users' satisfaction increased. Zhang et al., [20] presented a binary optimization method which incorporates the diversity of recommendation lists and the similarity between recommendation results and users' preferences. The experimental results show that this model outperformed the baselines in terms of diversity and accuracy. Zhou et al., [26] proposed a hybrid method to balance the accuracy and diversity of recommendations, which combines a diffusion-based recommendation algorithm and a "heat-spreading" algorithm addressing diversity problem. Furthermore, Shi et al., [27] argued that different users should have adaptive levels of diversity because of users' different interest extents and users' difference in terms of the estimated preference uncertainty, and then they proposed a new latent factor portfolio model which combines latent factor model and mean-variance analysis in portfolio theory.

In recent years, some novel approaches regarding recommendation diversity were proposed. Noia et al., [28] modeled users' propensity toward selecting diverse items, and then presented a diversification method to re-rank the list of top-N items predicted by a recommendation algorithm. Qin et al., [29] proposed a novel approach called contextual combinatorial bandit, where diverse items

[1] https://delicious.com/
[2] http://www.bibsonomy.org/
[3] https://www.flickr.com/