Knowledge-Based Systems 57 (2014) 57–68

Contents lists available at ScienceDirect

Knowledge-Based Systems

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

Merging trust in collaborative filtering to alleviate data sparsity and cold start

Guibing Guo *, Jie Zhang, Daniel Thalmann

School of Computer Engineering, Nanyang Technological University, Singapore

ARTICLE INFO

Article history:
Received 29 May 2013
Received in revised form 17 November 2013
Accepted 7 December 2013
Available online 12 December 2013

Keywords:
Recommender systems
Collaborative filtering
Cold start
Data sparsity
Trusted neighbors

ABSTRACT

Providing high quality recommendations is important for e-commerce systems to assist users in making effective selection decisions from a plethora of choices. Collaborative filtering is a widely accepted technique to generate recommendations based on the ratings of like-minded users. However, it suffers from several inherent issues such as *data sparsity* and *cold start*. To address these problems, we propose a novel method called "Merge" to incorporate social trust information (i.e., trusted neighbors explicitly specified by users) in providing recommendations. Specifically, ratings of a user's trusted neighbors are merged to complement and represent the preferences of the user and to find other users with similar preferences (i.e., similar users). In addition, the quality of merged ratings is measured by the confidence considering the number of ratings and the ratio of conflicts between positive and negative opinions. Further, the rating confidence is incorporated into the computation of user similarity. The prediction for a given item is generated by aggregating the ratings of similar users. Experimental results based on three real-world data sets demonstrate that our method outperforms other counterparts both in terms of accuracy and coverage.

1. Introduction

The emergence of Web 2.0 applications has greatly changed users' styles of online activities from searching and browsing to interacting and sharing [6,40]. The available choices grow up exponentially, and make it challenge for users to find useful information which is well-known as the *information overload* problem. Recommender systems are designed and heavily used in modern e-commerce applications to cope with this problem, i.e., to provide users with high quality, personalized recommendations, and to help them find items (e.g., books, movies, news, music, etc.) of interest from a plethora of available choices.

Collaborative filtering (CF) is one of the most well-known and commonly used techniques to generate recommendations [1,17]. The heuristic is that the items appreciated by those who have similar taste will also be in favor of by the active users (who desire recommendations). However, CF suffers from several inherent issues such as *data sparsity* and *cold start*. The former issue refers to the difficulty in finding sufficient and reliable similar users due to the fact that users in general only rate a small portion of items, while the latter refers to the dilemma that accurate recommendations are expected for the cold users who rate only a few items and thus whose preferences are hard to be inferred.

To resolve these issues and model user preferences more accurately, additional information from other sources is studied and incorporated into CF including friendship [19], membership [38,12] and social trust [41,2], where trust is believed less ambiguously and more reliable than friendship and membership. In this paper, trust is defined as one's belief toward others in providing accurate ratings relative to the preferences of the active user. Both implicit trust (e.g., [26,28]) and explicit trust (e.g., [4,8,25,27]) have been investigated in the literature. The former trust is inferred from user behaviors such as ratings whereas the latter is directly specified by users. By definition, the explicit trust tends to be more accurate and reliable than the implicit one. We focus on the explicit trust in this paper. Although many trust-based approaches have been proposed and the improvements to some extent have been achieved, there is still much room left for a better trust-based approach as stressed by [32].

In this paper, we propose a novel trust-based approach called "Merge" by incorporating the trusted neighbors explicitly specified by the active users in the systems, aiming to improve the overall performance of recommendations and to ameliorate the data sparsity and cold-start problems of CF. Specifically, we merge the ratings of trusted neighbors of an active user by averaging the ratings on the commonly rated items according to the extent to which the trusted neighbors are similar to the active user. The

* Corresponding author. Tel.: +65 84005322.
E-mail addresses: gguo1@ntu.edu.sg (G. Guo), zhangj@ntu.edu.sg (J. Zhang), danielthalmann@ntu.edu.sg (D. Thalmann).

0950-7051/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.knosys.2013.12.007

quality of the merged rating is measured by the confidence considering the number of ratings and the ratio of conflicts between positive and negative opinions (ratings). The set of merged ratings is then used to represent the active user's preferences and to find similar users based on user similarity. Further, the rating confidence is also taken into account in the computation of user similarity. Finally, the Merge method is incorporated into a conventional CF to generate recommendations. Experiments on three real-world data sets are conducted to demonstrate the effectiveness of our method in terms of accuracy and coverage. The results confirm that our method achieves promising recommendation performance, especially effective for the cold users comparing with the other counterparts. Although the idea of incorporating trust information into recommender systems is not new, our paper is the first work to effectively complement user rating profiles based on the ratings of trusted neighbors. Hence, our method shades light on a new way to build an effective trust-aware recommender system. A preliminary version of our work was published at the UMAP'12 conference [10].

The rest of this paper is organized as follows. Section 2 gives a brief overview of related research on trust-based CF from which the research gap is identified and motivating our present work. The proposed approach is then elaborated in Section 3 where we also highlight the advantages of our method in principle. Experiments on three real-world data sets are conducted in Section 4 to verify the effectiveness of our method in predicting items' ratings, especially for the cold users. Finally, Section 5 concludes our work and outlines potential future research.

2. Related work

Many CF approaches have been proposed in the literature to resolve the data sparsity and cold start problems. Generally, they can be classified into two categories: memory-based and model-based. The most well-known model is matrix factorization (MF) based approaches, such as SVD [20], NNMF [39], tensor factorization [31]. Model-based approaches usually can achieve better accuracy and coverage than memory-based approaches. This is because the former ones will train a prediction model using global rating data whereas the latter concentrate on local rating information. However, model-based approaches cannot properly explain how the recommendations are generated and effectively adopt new ratings due to trained static models. Although these problems might be mitigated to some extent such as [9,7], the real applications (e.g., citeulike.com and last.fm) usually adopt simpler memory-based approaches. Further, a lesson learned from the Netflix competition is that no single (memory- or modal-based) approach can always achieve the best performance, and different methods generally reveal different patterns of rating data [3]. Hence, it is necessary to further improve the performance of memory-based approaches.

To better model user preferences for the cold users who only rated a few items, additional information is often adopted. For example, [19] take into consideration both the social annotation (tag) and friendships inherently established among users in a music track recommender system. Due to the ambiguity of friendship, friends may have different preferences in items. In contrast, users joining the same online community are more likely to have similar preferences [38]. Hence by leveraging data from multiple channels including memberships in a project *wiki* [5,12] build a system named SONAR for people recommendation. Comparing with friendship and membership, trust information is of less ambiguity and more relevant to similarity [2,33,41]. Till now many trust-based approaches have been proposed [13,15,16,24], and trust has been demonstrated to be able to decrease recommendation errors and also increase recommendable items.

Trust information can be explicitly collected from users or implicitly inferred from users' rating information. The former trust is specified directly by users themselves. Typical applications are FilmTrust (trust.mindswap.org/FilmTrust/) and Epinions.com where each user can specify others as trustworthy or untrustworthy. In contrast, implicit trust is usually inferred from user behaviors, such as ratings. For example, [26] define the *profile-level* and *item-level* trust as the percentage of correct predictions from the views of general profile and specific items, respectively. [14,22] compute the implicit trust by the aggregation of value of a user's rating which is defined as the differences between the predicted ratings (based on only one user) and the ground truth. The intuition is that the closer two ratings are, the more value of that rating, and the more valuable ratings a user has, the more trustworthy the user will be. Since explicit trust is directly specified by users themselves, it is believed more accurate and reliable than implicit trust in determining the social relationships among users. In addition, [18] also show that letting users specify the explicit level of their trust (i.e., explicit trust) on the friends can improve the quality of recommendations. Hence, the present work focuses on the explicit trust.

Many approaches have been proposed to make use of the explicit trust. For example, [15] design the *TrustWalker* approach to randomly select trusted neighbors in the trust networks, where users are represented as nodes and trusted neighbors are connected with each other by trust links (i.e., edges) the strength of which indicates the trustworthiness between two users. Trust information of the selected neighbors is combined with an item-based technique to predict item ratings. In contrast, our work focuses on generating predictions by combining trust information with a user-based technique. [23] report that more accurate prediction algorithms are possible by incorporating trust information into traditional collaborative filtering. They do not directly use trust to substitute similarity but rather amplify similarity measurement by taking into account the number of messages exchanged among users. Hence, this approach is message specific. Further, a number of hybrid approaches incorporating trust are also proposed, such as [29,30]. Good performance can be achieved by combining both user- and item-based CF approaches. However, in this paper we focus on how to further improve the user-based CF using explicit trust.

The closest approaches to ours are as follows. [25] analyze the drawbacks of conventional CF-based recommender systems, and elaborate the rationale why incorporating trust can mitigate those problems. They propose the *MoleTrust* algorithm, which performs depth-first search, to propagate and infer trust in the trust networks. Empirical results show that the coverage is significantly enlarged but the accuracy remains comparable when propagating trust. Similarly, [8] propose a breadth-first search method called *TidalTrust* to infer and compute trust value. Both approaches substitute similarity with trust to predict item ratings, and the performance of the two algorithms is close [34]. Hence, we will only compare our method with one of them, namely MoleTrust in this paper. In addition, [4] propose to enhance CF by predicting the ratings of similar users who did not rate the target items according to the ratings of their trusted neighbors, so as to incorporate more similar users for recommendation. However, it performs badly in cold conditions where only few ratings are available, which is the main concern of the present work. Another recent work using the explicit trust network is proposed by [27]. They improve the prediction accuracy by reconstructing the trust networks. More specifically, the trust links between two users will be removed if their similarity is lower than a threshold. Empirical results show that good performance is achieved at the cost of poor coverage, and it fails to function in cold conditions where user similarity may not be computable.