Expert Systems with Applications 42 (2015) 4851–4858

Contents lists available at ScienceDirect

Expert Systems with Applications

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

Escaping your comfort zone: A graph-based recommender system for finding novel recommendations among relevant items

Kibeom Lee [a], Kyogu Lee [a,b,*]

[a] Music and Audio Research Group, Graduate School of Convergence Science and Technology, Seoul National University, Seoul, Republic of Korea
[b] Advanced Institutes of Convergence Technology, Suwon, Republic of Korea

A R T I C L E I N F O

Article history:
Available online 12 August 2014

Keywords:
Recommender systems
Novelty
Relevance
Popularity bias
Music recommendation

A B S T R A C T

Recommender systems have steadily advanced in their ability to filter out unnecessary information and deliver the most relevant data to users. Such recommender systems are being used commercially with popular methods being based on collaborative filtering. While collaborative filtering-based recommenders perform well in terms of accuracy, they lack the ability of finding fresh and novel items, due to the nature of its inner workings. We propose a new graph-based recommender system that uses only positively rated items in users' profiles to construct a highly-connected, undirected graph, with items as nodes and positive correlations as edges. Using the concept of entropy and the linked items in the graph, the proposed system can find recommendations that are both novel and relevant. We test the system on Last.fm data to recommend music to users and show that the proposed recommender system is indeed able to provide novel recommendations while keeping them relevant to the user profile, consistently outperforming a state-of-the-art matrix factorization-based recommender.

1. Introduction

Recommender systems have come a long way ever since the introduction of collaborative filtering (Resnick, Iacovou, Suchak, Bergstrom, & Riedl, 1994; Shardanand & Maes, 1995), a method of predicting users' interests by aggregating the preferences or opinions of other similar users. Its key concept came from our everyday lives, in which we share our opinions that spread through word-of-mouth. These opinions help others make decisions, which in this case, receive recommendations.

Collaborative filtering has become an integral part of e-commerce by tackling the problem of information overload, which rose due to advances in the Internet and drop in hardware costs. The first commercial recommender to use collaborative filtering was Tapestry (Goldberg, Nichols, Oki, & Terry, 1992), while the most representative recommender system used successfully today would be that of Amazon's (Linden, Smith, & York, 2003).

Although collaborative filtering has steadily advanced through the years with improved prediction accuracy and its popularity in the commercial scene, it suffers from popularity bias, an inherent problem that stems from collaborative filtering's social component. This leads to recommendations composed of popular items, which can be accurate in terms of relevancy but also highly probable of being familiar to the user. While accurate recommendations might be correct recommendations, they may not be valuable or useful recommendations for the user.

Incidentally, the literature on recommender system metrics have also realized that there is more to it than well-known, accurate recommendations regarding recommender systems (McNee, Riedl, & Konstan, 2006). Thus, more recent studies have expanded the focus of recommendations from accuracy to other factors, such as diversity and novelty (Herlocker, Konstan, Terveen, & Riedl, 2004).

Thus, there is a need to find a way to find recommendations composed of long-tail (Anderson, 2006) items that are unknown to the user while being relevant. In this paper, we propose a novel approach to finding lesser-known and relevant recommendations, using an item graph constructed from positively co-rated items. This graph is used to extract a sub-graph representing the items of a user's profile their and neighboring items, which is then analyzed using the concept of Shannon's entropy to find novel and relevant recommendations. The system's overview is illustrated in Fig. 1.

2. Related work

Despite the rising attention regarding the need of novelty and diversity-focused recommender systems there have only been

* Corresponding author at: Music and Audio Research Group, Graduate School of Convergence Science and Technology, Seoul National University, Seoul, Republic of Korea.
E-mail address: kglee@snu.ac.kr (K. Lee).

http://dx.doi.org/10.1016/j.eswa.2014.07.024
0957-4174/

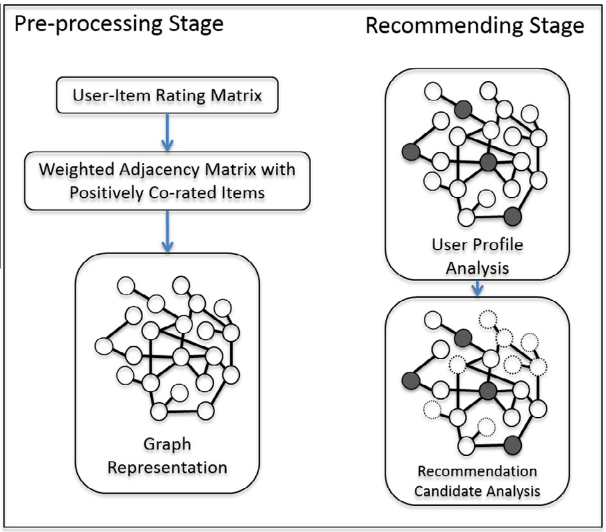

Fig. 1. Overview of proposed system.

few studies addressing novelty-focused recommender systems (Domingues et al., 2013; Lee & Lee, 2011, 2013, 2014; Lee, Yeo, & Lee, 2010; Park & Tuzhilin, 2008), whereas a large portion of studies focused on diversity. Diversity-focused recommenders such as Adomavicius and Kwon (2011), Chandrashekhar and Bhasker (2011), Gan and Jiang (2013), Lee, Park, Kahng, and Lee (2013) used various methods such as network flow maximization, selective predictability and entropy, network-based collaborative filtering, and random walk-based node ranking to increase the diversity of the recommendations.

However, finding diverse items and finding novel items are two different tasks (Abbassi, Amer-Yahia, Lakshmanan, Vassilvitskii, & Yu, 2009). A recommendation list can be diverse yet not novel, e.g. a movie recommendation list composed of popular items from various genres. While these systems were aimed to find diverse recommendations, the purpose of our study is to search for novel recommendations that are also relevant. This is an important contribution for the recommender system community as research on methods of minimizing the trade-off between novelty and relevancy is not as active compared to other factors such as diversity.

Regarding novelty-oriented recommenders, Iaquinta et al. (2008) presented a content-based recommender system combined with serendipitous heuristics to deliver novel recommendations. The system used text documents that contained descriptions of the items to analyze the content and measure similarity and serendipity. Experiments were carried out from a corpus of 45 paintings with accompanying metadata on 30 users. A significant drawback of this system is that, as with all content-based recommenders, it needs additional metadata to find recommendations. Also, the dataset used in the experiments is significantly smaller than those typically used from available datasets. In this study, we propose a recommender system that only needs user-item rating data to provide both novel and relevant recommendations. We also test our system with a large dataset consisting of approximately 185,000 items and 3 million ratings.

Abbassi et al. (2009) proposed an Outside-The-Box (OTB) recommender system that found fresh discoveries while maintaining relevancy. It used clusters of items, called item regions, to find underexposed items for recommendation. Experiments showed that the OTB recommendations contributed to over 80% of the overall nDCG (Normalized Discounted Cumulative Gain), indicating that the recommendations were novel and potentially useful. However, this study does not continue to verify whether these potentially useful recommendations were *actually* useful as the relevance of truly novel recommendations can only be verified through explicit user feedback (i.e. live user tests). In this paper, we perform a live user test to receive user feedback on the relevance of the novel recommendations provided by the proposed system and compare the results with a state-of-the-art matrix factorization-based recommender.

Lee and Lee (2014) introduced a recommender system that used the concept of experts among the user population to generate novel and relevant recommendations. The items were grouped in predefined clusters, in which users could be Experts depending on their profile. Experiments with widely used metrics showed that the system outperformed a matrix factorization-based algorithm in terms of finding novel items with only a slight trade-off with accuracy. The recommender system performs well in finding novel and relevant items, but has too many variables that have to be found heuristically. In addition, values that may work for a certain dataset may not work well for a different dataset. Thus, there is a great overhead in adjusting the values of the many parameters the algorithm offers. The study also did not perform live user tests for measuring the relevance of novel recommendations. In this paper, we propose a much simpler recommender system that requires a single parameter to adjust the novelty and perform a live user test for evaluation.

3. Positively-linked graph-based recommender

In this section, we present a graph-based recommender system (GraphRec) targeted to provide novel and diverse recommendations. The algorithm can be divided into three large steps. First, the rating matrix, or user-item matrix, is converted to a weighted adjacency matrix. Next, this newly created matrix is used to create a graph of items with links representing positive relations. With this graph, we calculate the entropy for all nodes, completing the pre-processing stage.

3.1. Weighted adjacency matrix

The first step of the proposed system is to compose a weighted adjacency matrix to create the graph. Here, we assume that there is a user-item matrix, or rating matrix at hand, with dimensions $N * M$, indicating N users and M items. We create an $M * M$ adjacency matrix and begin filling it out.

To populate the adjacency matrix, we find and group the positively-related items together. Positively-related items are defined as the set of items in a user's profile that have ratings higher than the user's average rating. Formally, the set of positively-related items, P, would be defined as in Eq. (1), where i indicates an item in the user profile U; U_i is the rating of Item i in the user's profile U; and $\overline{U}$ is the mean rating of the user's profile.

$$P = \{i | U_i \geqslant \overline{U}, \quad i \in U\} \tag{1}$$

For each user, we find the positively-related items and increment every possible combination of paired items in the adjacency matrix by 1.

The following example illustrates this process. For this purpose, we assume that there are 5 items $\{A, B, C, D, E\}$ and a single user U. User U's profile has the following ratings for the items previously mentioned in their respective orders: $\{5, 3, 3, 2, 1\}$. The mean profile rating, $\overline{U} = 2.8$. Thus, the positively-related item set $P = \{A, B, C\}$. With P, we create the following weighted adjacency matrix (left). When processed for all users, we obtain a weighted adjacency matrix, M, showing items with positive relations. We populate the matrix randomly for illustration purposes (right).