Available online at www.sciencedirect.com
ScienceDirect
Information Processing and Management 43 (2007) 473–487
INFORMATION PROCESSING & MANAGEMENT
www.elsevier.com/locate/infoproman

A probabilistic music recommender considering user opinions and audio features

Qing Li [a,*], Sung Hyon Myaeng [a], Byeong Man Kim [b]

[a] *Information and Communications University, Daejeon, 305-732, Republic of Korea*
[b] *Kumoh National Institute of Technology, Kumi, 730-701, Republic of Korea*

Received 27 May 2006; accepted 25 July 2006
Available online 9 October 2006

Abstract

A recommender system has an obvious appeal in an environment where the amount of on-line information vastly outstrips any individual's capability to survey. Music recommendation is considered a popular application area. In order to make personalized recommendations, many collaborative music recommender systems (CMRS) focus on capturing precise similarities among users or items based on user historical ratings. Despite the valuable information from audio features of music itself, however, few studies have investigated how to utilize information extracted directly from music for personalized recommendation in CMRS. In this paper, we describe a CMRS based on our proposed item-based probabilistic model, where items are classified into groups and predictions are made for users considering the Gaussian distribution of user ratings. In addition, this model has been extended for improved recommendation performance by utilizing audio features that help alleviate three well-known problems associated with data sparseness in collaborative recommender systems: user bias, non-association, and cold start problems in capturing accurate similarities among items. Experimental results based on two real-world data sets lead us to believe that content information is crucial in achieving better personalized recommendation beyond user ratings. We further show how primitive audio features can be combined into aggregate features for the proposed CRMS and analyze their influences on recommendation performance. Although this model was developed originally for music collaborative recommendation based on audio features, our experiment with the movie data set demonstrates that it can be applied to other domains.

Keywords: Collaborative filtering; Music recommender system; Probabilistic model; Information filtering

1. Introduction

With the development of e-commerce and the proliferation of easily accessible information, recommender systems have become a popular technique to prune large information spaces so that users are directed toward

* Corresponding author.
E-mail address: liqing@icu.ac.kr (Q. Li).

0306-4573/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ipm.2006.07.005

those items that best meet their needs and preferences. The technique has been an integral part of e-commerce sites such as Amazon, Yahoo and CDNow.

At the initial stage, recommender systems mostly relied on a content-based filtering (CBF) mechanism. It selects the right information for a user by matching the user preference, which may be implicit or explicit, against databases. For example, a search engine recommends web pages whose contents are similar to a user query (Wittenburg, Das, Hill, & Stead, 1995). Despite the efficiency of a CBF in locating textual items relevant to a topic, it is a little difficult for such a system to be directly applied to multimedia when they lack textual descriptions such as a music genre, title, composer, or singer. In addition, representing a user preference or a query for non-textual features is also a big challenge for a music recommender system (MRS). Some researchers (Ghias, Logan, Chamberlin, & Smith, 1995) suggested a method of querying an audio database by humming. Not all users, however, have a gift of humming the melody of their favorite songs for searching.

Collaborative filtering (CF), as in GroupLens (Resnick, Iacovou, Suchak, Bergstorm, & Riedl, 1994) and Ringo (Shardanand & Maes, 1995), has been considered a mainstream technique for recommender systems for a long time until now. CF uses opinions of others to predict the interest of a user. A target user is matched against the database of user histories to discover other users, called neighbors, who have historically similar tastes. Items the neighbors like are then recommended to the target user. For instance, the GAB system (Wittenburg et al., 1995) recommends web pages based on the bookmarks; the Jeter system recommends jokes (Gupta, Dhruv, Mark Digiovanni, Hiro Narita, & Jester, 1999); MovieLens recommends movies (http://movielens.umn.edu); and Flycasting recommends online radio programs (Hauver, 2001). Most of prevalent CF systems focus on calculating the user–user similarity to make predictions, which is the so-called user-based CF. However, Sarwar (Sarwar, Karypis, Konstan, & Riedl, 2001) has proven that an item-based CF is better than a user-based CF on precision and computational complexity.

Despite the popularity of CF techniques, researchers also realized that the content information of items did help providing a good recommendation service. The idea of a hybrid system capitalizing on both CF and CBF has been suggested. Examples include Fab (Balabanovic & Shoham, 1997), ProfBuilder (Wasfi, 1999), RAAP (Delgado, Ishii, & Ura, 1998), and Chen's music recommender (Chen & Hung-Chen, 2001).

The Fab system (Balabanovic & Shoham, 1997) uses content-based techniques instead of user ratings to create user profiles. Since the quality of predictions is fully dependent on the content, inaccurate profiles result in inaccurate correlations with other users, yielding poor predictions. ProfBuilder (Wasfi, 1999) recommends web pages using both content-based and collaborative filters, each of which creates a recommendation list separately to make a combined prediction. Claypool et al. (1999) apply a hybrid approach for an online newspaper domain, combining two types of predictions using an adaptive weighted average. As the number of users accessing an item increases, the weight of the collaborative component tends to increase. However, the authors do not clearly describe how to decide the weights of collaborative and content-based components. RAAP (Delgado et al., 1998) is a content-based collaborative information filtering system that helps the user classify domain specific information found on the WWW, and also recommends those URLs to other users with similar interests. To determine the similarity of interests among users, a scalable Pearson correlation algorithm based on the web page category is used.

In Chen's music recommender system (Chen & Hung-Chen, 2001), both user interests and music features are applied for making recommendations. Access histories of users are analyzed to mine user interests. This work makes a content-based and collaborative recommendation, emphasizing on the degree to which a user favors a music group. This approach assumes that each piece of music is assigned to a group. If a user is interested in a certain group, the songs in the group are recommended to the user with an equal value. In other words, this system can only recommend an un-ordered list of music pieces to users.

Clustering is the key idea of CF including hybrid systems that rely on CF. Clustering is implemented in various ways, either explicitly or implicitly in most CF techniques. A pioneering CF method (Resnick et al., 1994) based on Pearson correlation coefficient aims at grouping the users with similar tastes for an active user. It first calculates the similarities between the active user and others, and then clusters users into groups with similar tastes. There are several variants (Breese, Heckerman, & Kardie, 1998; Li, Qing, Kim, Guan, & Oh, 2004; Sarwar et al., 2001) derived from the CF method. Since relationships should be calculated and saved in the memory for recommendation, these methods are called as memory-based techniques.