Expert Systems With Applications 55 (2016) 546–558

Contents lists available at ScienceDirect

Expert Systems With Applications

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

Online recommender system for radio station hosting based on information fusion and adaptive tag-aware profiling

Dmitry I. Ignatov[a,*], Sergey I. Nikolenko[b,c], Taimuraz Abaev[a], Jonas Poelmans[a]

[a] Computer Science Faculty, National Research University Higher School of Economics, Kochnovskiy proezd, 3, Moscow 125319, Russia
[b] Laboratory of Internet Studies, National Research University Higher School of Economics, ul. Soyuza Pechatnikov, d. 16, St. Petersburg 190008, Russia
[c] Steklov Mathematical Institute, nab. r. Fontanka, 27, St. Petersburg 191023, Russia

ARTICLE INFO

Article history:
Received 22 July 2014
Revised 13 February 2016
Accepted 14 February 2016
Available online 18 February 2016

Keywords:
Music recommender system
Interactive radio network
Hybrid recommender system
Information fusion
Adaptive tag-aware profiling
Implicit feedback

ABSTRACT

We present a new recommender system developed for the Russian interactive radio network *FMhost*. To the best of our knowledge, it is the first model and associated case study for recommending radio stations hosted by real DJs rather than automatically built streamed playlists. To address such problems as cold start, gray sheep, boosting of rankings, preference and repertoire dynamics, and absence of explicit feedback, the underlying model combines a collaborative user-based approach with personalized information from tags of listened tracks in order to match user and radio station profiles. This is made possible with adaptive tag-aware profiling that follows an online learning strategy based on user history. We compare the proposed algorithms with singular value decomposition (SVD) in terms of precision, recall, and normalized discounted cumulative gain (NDCG) measures; experiments show that in our case the fusion-based approach demonstrates the best results. In addition, we give a theoretical analysis of some useful properties of fusion-based linear combination methods in terms of graded ordered sets.

1. Introduction

Music recommendation is an important and challenging direction in the field of recommender systems. It is a hard problem to extract relevant information for similarity search from large music collections, and sources of rich semantic metadata are often needed (Celma, 2010). Many recent works in this area have appeared at the International Society for Music Information Retrieval Conference (ISMIR) (Müller & Wiering, 2015) and the Recommender Systems conference (RecSys) (Werthner, Zanker, Golbeck, & Semeraro, 2015).

Recently, the focus of computer science research in music information studies has shifted from pure music information retrieval and exploration (Gleich, Zhukov, Rasmussen, & Lang, 2005; Hilliges, Holzer, Klüber, & Butz, 2006) to music recommendations (Brandenburg et al., 2009; Celma, 2010). It is not a new direction (Avesani, Massa, Nori, & Susi, 2002); however, it is now inspired by new capabilities of large online services that can provide not only millions of tracks for listening but also thousands of radio stations to choose from on a single web site. Moreover, social tagging is an important factor that paves the way to new recommender algorithms based on tag similarity (Nanopoulos, Rafailidis, Symeonidis, & Manolopoulos, 2010; Symeonidis, Ruxanda, Nanopoulos, & Manolopoulos, 2008b; Yang, Bogdanov, Herrera, & Sordo, 2012).

Despite many high quality works on different aspects of music recommendation, there are only a few studies devoted to online radio station recommender systems (Aizenberg, Koren, & Somekh, 2012; Grant, Ekanayake, & Turnbull, 2013). Several radio-like online broadcasting services, including *last.fm, Yahoo!LaunchCast*, and *Pandora*, are known for their recommender systems and work on a commercial basis (however, the latter two do not operate in Russia). There is their Russian counterpart, Yandex.Radio[1]; however, as the aforementioned services, it only provides access to radio stations with playlists, automatically composed according to a catalogue of selection criteria. Recommendation of real online radio stations is different. It presents difficulties not only due to the musical content but also because a recommender needs to find relevant dynamically changing objects while usually relying on implicit feedback only.

Thus, in this work we consider the music recommendation problem from a slightly different angle. We consider the Russian online radio hosting service *FMhost*, in particular, its new hybrid recommender system. First, we recommend radio stations, i.e.,

* Corresponding author. Tel.: +7 9263818033.
E-mail addresses: dignatov@hse.ru, dmitrii.ignatov@gmail.com (D.I. Ignatov), sergey@logic.pdmi.ras.ru (S.I. Nikolenko).

[1] http://radio.yandex.ru.

http://dx.doi.org/10.1016/j.eswa.2016.02.020
0957-4174/

sequences or sets of compositions, which are manually composed by real DJs and dynamically change, rather than individual tracks as most other music recommenders do. Second, as we demonstrate below, the *FMhost* service does not have enough data for reasonable SVD-based recommendations; still, recommendations have to be provided. To overcome these problems, we propose a novel recommender algorithm that combines three data sources: radio station visits, listening events for music with specific tags, and frequency of tags applied to radio stations and their content. We show experimental results on the *FMhost* dataset and propose a fusion of two different algorithms that can be tuned for specific quality metrics, e.g., precision, recall, and NDCG.

The paper is organized as follows. In Section 2, we briefly survey related work in music recommendation. Section 3 outlines the online radio service *FMhost*. In Section 4, we propose a novel recommender model, two basic recommender algorithms, a third algorithm that combines them, and describe the recommender system architecture. Section 5 provides examples, and Section 6 discusses the basic principles and problems in details. Quality of service (QoS) measurements, a comparison with an SVD-based approach, and certain insights into *FMhost* user behaviour are discussed in Section 7. In Section 8, we provide a theoretical justification of the chosen aggregation of rankings. Section 9 concludes the paper.

2. Related work

Music recommendation becomes especially important because modern systems that provide music to their users aim to take into account infrequently requested musical compositions and/or collections of compositions such as radio stations from the long tail of the distribution. Most music recommender systems work under the general principles of collaborative filtering (Koren & Bell, 2011). For instance, *last.fm* mines user tastes both explicitly, from likes with which the users mark compositions, and implicitly, from compositions that the users actually listen to. Many music recommenders also incorporate content-based features, analysing the actual composition. An interesting example of the latter is the *Pandora* service (Castelluccio, 2006; Joyce, 2006). *Pandora* decomposes a music composition into the so-called music genome; the "genes" of a composition are different for different music genres, and compositions are graded for various genes by professional musicologists. Still, even content-based systems usually employ collaborative filtering and use content features to supplement classical recommender algorithms, so basic ideas of collaborative filtering such as matrix factorization (Koren, 2008; Koren, Bell, & Volinsky, 2009; Ma, Yang, King, & Lyu, 2009; Salakhutdinov & Mnih, 2008a; 2008b) and nearest neighbors in both user-based and item-based collaborative filtering (Koren, 2010; Resnick, Iacovou, Sushak, Bergstrom, & Riedl, 1994; Said, Jain, Kille, & Albayrak, 2012; Sarwar, Karypis, Konstan, & Riedl, 2001) certainly apply; see also recent surveys on recommender systems (Bobadilla, Ortega, Hernando, & Gutiérrez, 2013; Konstan & Riedl, 2012; Lü et al., 2012; Zhou, Xu, Li, Josang, & Cox, 2012).

A widely acclaimed public contest on music recommender algorithms, KDD Cup,[2] was recently held by *Yahoo!*. In KDD Cup, track 1 was devoted to learning to predict users' ratings of musical items (tracks, albums, artists, and genres) where the items formed a taxonomy: tracks belong to albums, albums belong to artists, where albums and tracks are also tagged with genres. Track 2 aimed at developing learning algorithms for separating music tracks scored highly by specific users from tracks that have not been scored by them. It attracted a lot of attention to the problems that are both typical for recommender systems and specific for music recommendation: scalability issues, capturing the dynamics and taxonomic properties of the items (Koenigstein, Dror, & Koren, 2011). Another major music recommender contest with open data, the Million Songs Dataset Challenge,[3] was held in 2012 by the Computer Audition Lab at UC San Diego and LabROSA at Columbia University (McFee, Bertin-Mahieux, Ellis, & Lanckriet, 2012). The core data consists of triples (*user, song, count*); it covers approximately 1.2 million users and more than 380,000 songs. While music recommender problems seldom have explicit user preferences (it is unlikely that users would rate every track they listen to), this kind of data can serve as input to one-class recommender algorithms that operate on only one kind of activity as input (Pan & Scholz, 2009; Pan et al., 2008; Sindhwani, Bucak, Hu, & Mojsilovic, 2010); these algorithms can also benefit from additional information about both users and items (Li, Hu, Zhai, & Chen, 2010; Yuan, Cheng, Zhang, Liu, & Lu, 2013).

Recent music recommendation trends reflect the advantages of hybrid approaches and call for user-centric quality measures (Celma & Lamere, 2011). For instance, the work (Hu & Ogihara, 2011) proposes a novel approach based on the so-called "forgetting curve" to evaluate "freshness" of predictions. Tags (both moderated and user-generated) turn out to be especially important for recommendations. The work (Bogdanov & Herrera, 2011) studies the problem of how much metadata one needs in music recommendation: a subjective evaluation of 19 users has revealed that pure content-based methods can be drastically improved with genre tags. Finally, the authors proposed a recommender approach that starts from an explicit set of music tracks provided by the user as evidence of his/her preferences and then infers high-level semantic descriptors for each track (Bogdanov et al., 2013). There is also a rich body of work devoted to tag suggestion and automated tag generation for music compositions since tags, especially generic tags like genres or social tags like "music for jogging" can then serve as preconstructed recommenders for the users (Bergstra, Casagrande, Erhan, Eck, & Kégl, 2006; Eck, Lamere, Bertin-Mahieux, & Green, 2007; Font, Serrà, & Serra, 2012; Zhao et al., 2010). However, in their extensive survey on music recommendation and retrieval (Kaminskas & Ricci, 2012) wrote "unfortunately, research on semantic music retrieval is still in its early stages," and call to use tag-based approaches since "pure content analysis often fails to capture important aspects of human perception of music tracks." Thus, in Hyung, Lee, and Lee (2014) the authors extracts keyterms from user's personal stories and match their profiles with songs by latent semantic analysis. From graph-theoretic point of view, a quite close work is done by Lee, Kahng, and Lee (2015). Thus, the authors use so called aggregated bipartite factor-items graphs where factors should eliminate the direct usage of heterogeneous metadata reducing the number of stored graph nodes. In fact, our tag-aware profiles aggregate listened tracks considered as graph nodes into tag-based factors. However, we do it dynamically. Existing tag-aware frameworks are mostly used for building static profiles or tag recommendations per se (Kim & Kim, 2014; Nanopoulos et al., 2010); they also use different tag-propagation mechanism to cope with sparsity. In our solution we use online tag propagation for dynamic updating of users and radios profiles.

A substantial part of works on music recommendation is devoted to personalised playlist generation (Balkema & van der Heijden, 2010; Liu, Hsieh, & Tsai, 2010; Mocholí, Martinez, Martínez, & Catalá, 2012). And many online services (e.g., *last.fm* and *LaunchCast*) call their audio streams "radio stations", however they are actually playlists from a database of tracks which are based on a recommender system rather than real predefined channels.

[2] http://kddcup.yahoo.com/.

[3] http://labrosa.ee.columbia.edu/millionsong/.