Artificial Intelligence 219 (2015) 25–39

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Learning pseudo-tags to augment sparse tagging in hybrid music recommender systems

Ben Horsburgh, Susan Craw *, Stewart Massie

School of Computing Science & Digital Media and IDEAS Research Institute, Robert Gordon University, Aberdeen, UK

ARTICLE INFO

Article history:
Received 25 March 2013
Received in revised form 12 November 2014
Accepted 18 November 2014
Available online 25 November 2014

Keywords:
Music recommendation
Hybrid representations

ABSTRACT

Online recommender systems are an important tool that people use to find new music. To generate recommendations, many systems rely on tag representations of music. Such systems, however, suffer from tag sparsity, whereby tracks lack a strong tag representation. Current state-of-the-art techniques that reduce this sparsity problem create hybrid systems using multiple representations, for example both content and tags. In this paper we present a novel hybrid representation that augments sparse tag representations without introducing content directly. Our hybrid representation integrates pseudo-tags learned from content into the tag representation of a track, and a dynamic weighting scheme limits the number of pseudo-tags that are allowed to contribute. Experiments demonstrate that this method allows tags to remain dominant when they provide a strong representation, and pseudo-tags to take over when tags are sparse. We show that our approach significantly improves recommendation quality not only for queries with a sparse tag representation but also those that are well-tagged. Our hybrid approach has potential to be extended to other music representations that are used for recommendation but suffer from data sparsity, such as user profiles.

1. Introduction

During the past decade a massive shift has occurred in the way that people store, listen to and discover music. Traditionally, either a single or album was stored on external media such as cassette or CD. With advances in hardware and the growth of the world wide web, this storage has moved to large music collections on personal hard-drives and websites offering massive catalogues of music.

An interesting new problem that has arisen, as a consequence of this shift, is how to find and discover new music. Radio stations and music magazines continue to offer advice but, with such a large amount of music readily available, many people now look to recommender systems provided by online services. Users provide a query, such as an example track, and a recommender system will then return a set of recommendations to the user.

Core to this task of providing recommendations is the representation of tracks. Each track in the collection must have a corresponding representation, which describes the key facets of the track. Many state-of-the-art music recommender systems make use of tag-based representations, for example the hugely popular Last.fm service [1]. The tags typically consist of two components: standardised meta-data, such as artist, track title and genre; and free-text that users provide. Free-text tags are extremely useful to a recommender system, providing a wealth of different information, such as genres, artists,

* Corresponding author at: IDEAS Research Institute, Robert Gordon University, Garthdee Road, Aberdeen AB10 7GJ, UK. Tel.: +44 1224 262711.
E-mail addresses: s.craw@rgu.ac.uk (S. Craw), s.massie@rgu.ac.uk (S. Massie).

http://dx.doi.org/10.1016/j.artint.2014.11.004
0004-3702/

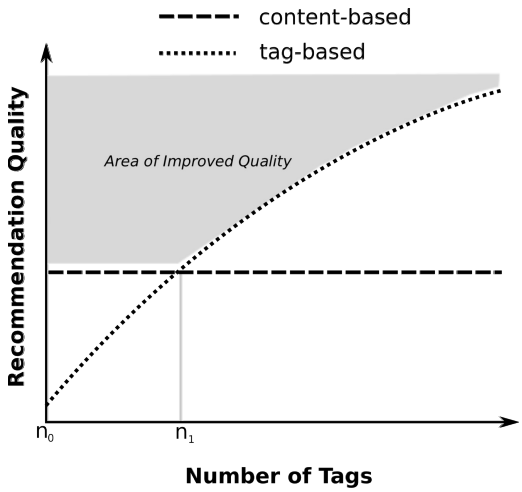

Fig. 1. The cold-start problem.

topics, opinions and contexts [2,3]. Many of these tags provide useful information for recommendation, for example, tags relating to genres, artists and topics. Context such as location can be used to improve personalised recommendations, such as when it is known a user is working in a given location [4,5].

Uninformative tags can be problematic for tag-based recommender systems, but a far larger problem is when very few tags exist at all. This is a problem commonly known as cold-start and arises from a large popularity bias in music listening habits [6]. As a result, many tracks are under tagged; a track is tagged if it is listened to, but a track is only recommended, and therefore listened to, if it already has tags.

Fig. 1 displays the general pattern of results found in the experiments, and here is used to illustrate the cold-start problem in query-by-example tag-based recommender systems. The horizontal axis shows the number of tags a query has, and the vertical axis shows recommendation quality. The dotted curve in Fig. 1 represents average recommendation quality in a tag-based recommender system. As the number of tags increases, so does the quality of tag-based recommendation, since a larger amount of meaningful knowledge is available. The horizontal dashed line in Fig. 1 represents average recommendation quality in a content-based recommender system. Since this system does not rely on tags, the quality of recommendations does not vary as the number of tags changes. For a low number of tags, content-based recommendations are better than tag-based, but as tags increase the tag-based system becomes better than content.

A simple hybrid approach, which can reduce the effects of the cold-start problem, is to use content-based methods for low tagged queries, and then switch to tag-based methods. This hybrid does not offer an improvement over any individual method, but will produce better recommendations as an overall system. In this simple hybrid approach, however, deciding when to change from content- to tag-based recommendation can cause problems.

Ideally, the switch from content- to tag-based recommendation is made at point n_1 in Fig. 1. In this case, the maximum performance is achieved for all numbers of tags. However, in practice, it is very difficult to get this switch over correct. If the switch is made too early then performance of the system is decreased by not using the strongest representation. Similarly, if the switch is made too late, then the full potential of the tag-based recommender is not taken advantage of. Further, the extreme case of cold-start is at point n_0, where a query track has no tag representation. This does not affect content-based recommendation, but will reduce a tag-based system to making random recommendations.

The aim of a hybrid system designed to reduce cold-start is to consistently provide recommendations in the shaded area. In this case, the hybrid system is able to take advantage of both content and tags, and through their combination improve recommendation quality in all cases. Making recommendations in this area is the motivation for developing our approach to cold-start recommendation using a hybrid recommender system.

In this paper related work on the cold-start problem, hybrid recommender systems, and their constituent parts is discussed first. We then define and present our approach to learning pseudo-tags from content, which guarantees a tag-based representation for every track in a collection. Next, we present our novel hybrid representation, which is a dynamic combination of both pseudo-tags and tags. To evaluate our approach, we then examine both the recommendation quality, and the quality of pseudo-tagging. Finally, we present our results and conclusions.

2. Related work

Tag-based representations, and the knowledge they describe, are very important to music recommender systems. For certain recommendation problems, tags are able to offer greater performance than either content-based or collaborative filtering methods [7]. Whilst a lot of progress has been made in how this knowledge can best be used, cold-start continues to cause problems in many tag-based systems [2].