

Can a Rock song have a Jazz audience? Relationship between folksonomy and collaborative filtering in music recommender systems.

Eugenio Tacchini
Università degli Studi di Milano
Via Bramante, 65
26013, Crema (CR), Italy
eugenio.tacchini@unimi.it

Ernesto Damiani
Università degli Studi di Milano
Via Bramante, 65
26013, Crema (CR), Italy
ernesto.damiani@unimi.it

ABSTRACT

In this paper we investigate the relationship between a folksonomy-based music classification and a music classification based on collaborative filtering, i.e. on the users' listening behavior. We found a correlation between folksonomy-based songs clustering and clustering computed using methods based on the audience listening behaviour and, using a combination of the two approaches, we also computed the eclecticism level of a sample set of users, finding that eclecticism seems to be a characteristic which changes according to the genre of music most loved by a user.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval – *Clustering, Information filtering.*

General Terms

Algorithms, Measurement, Human Factors.

Keywords

Collaborative filtering, recommender systems, music, eclecticism, clustering, folksonomy, music genres.

1. INTRODUCTION

Music folksonomy has become a very important part of the social music services ecosystem. Thanks to folksonomy, final users can describe artists and songs exactly as they perceive them (regardless of a priori classifications made by experts). Folksonomy tags have been used, alone or together with other data, as content metadata for recommender systems [1, 2].

In particular, music recommendation services often rely on folksonomy tags as a description of an artist/song as seen through the eyes of a community and such a description can be exploited to provide better recommendations.

In this paper we investigate the relationship between a folksonomy-based music classification and a music classification based on collaborative filtering, i.e. on the users' listening behavior. Our first research question can be stated as follows: "Does folksonomy-based songs clustering compute the same clusters as methods based on the audience listening behaviour?". Getting a clear idea of the relationship between folksonomy-based

and collaborative filtering music classification is important not only to understand the dynamics behind social music communities but also to help in the process of combining the two classification components in the context of hybrid recommender systems.

We then used a combination of the two classification methods analyzed in order to compute the eclecticism level of a sample set of users (i.e. the extent to which users listen to songs apparently far from their taste), coming to a second research question: "Is the eclecticism level of a user dependent on the music mostly loved by that user?", to which we tried to answer in the last part of this paper.

2. RELATED WORK

Music folksonomies has already been studied in the past: an example can be found in [3], where the agreement between expert-based music genres vocabularies and folksonomies was analyzed, coming to the conclusion that, while for some genres experts and wisdom of crowds agree, for other genres clearly disagree.

If we consider our first research question, the closest work is [4], where the authors investigated whether the tagging of artists is coherent with the artist similarities found with collaborative filtering techniques; however, in order to measure the artists similarities, they relied on the "similar artists" feature of last.fm: while the feature algorithm is not public, it is likely that it, at least nowadays (years ago probably the algorithm was different), it is not a pure collaborative filtering approach and the folksonomy itself is exploited in order to provide artists similarities. We instead calculate the collaborative-filtering similarities at the song level (not at the artist level) and we use the well-known k-means clustering algorithm.

To the best of our knowledge, there hasn't been any previous computer science work that addresses our second research question.

3. METHODOLOGY

3.1 Dataset used

The dataset used for the experiments contains the whole listening history of 992 unique users of the social radio last.fm¹. The

¹ <http://www.last.fm>

clusters in order to understand if there were relationships with the music genres. The data are reported in table 1.

Table 1. Composition of the resulting clusters

Cluster	Songs count	Genres composition	
		Dominant genre	Other genres
1	4,487	metal 88%	classic rock 9% hip-hop 1% pop 1%
2	4,748	classic rock 69%	pop 16% jazz 10% hip-hop 4% metal 1%
3	4,529	hip-hop 86%	jazz 8% classic rock 5% pop 2%
4	2,773	jazz 80%	classic rock 13% hip-hop 5% pop 2%
5	3,223	pop 86%	hip-hop 10% jazz 3% classic rock 1%

The resulting clusters are quite balanced in terms of the total number of songs, and the data clearly shows that, for each cluster, there is a dominant genre: in four cases over five, at least the 80% of the songs belong to the same genre. Since our clustering is based on the the distance among songs calculated according to the users listening behavior, the approach can be considered very similar to the first step usually processed by a classical item-based collaborative filtering algorithm in the recommender systems domain; this suggests that there should be a strong relationship between the results of an item-based collaborative filtering approach and the results of a content-based (folksonomy) approach.

Analyzing the genre composition, we can also observe, for each genre, which are the “closest” genres. While most of the results are expected, e.g. we already suspected that heavy metal songs have something in common with classic rock songs and not much with pop songs, some others are quite surprising: hip-hop and jazz seem to be not so far from each others.

4.2 Users' eclecticism level

Our second research questions was about the eclecticism level of users: do most of the users play songs belonging to the same cluster or to different clusters? Which are the most eclectic category of users? To answer this question, we computed, for each user, the composition (in term of clusters frequently played) of his listening history.

In order to considered a cluster as “played” by a user, we decided to set a threshold (k) on the percentage of songs played by that user and belonging to that cluster respect to the total number of

songs played by the same user. The reason is that having listened just a few songs from a cluster doesn't mean that a user likes that particular kind of songs: the plays could be “casual”, especially if we consider that the listening history provided by last.fm includes both the songs directly played by the users (e.g. from iTunes) and the songs played by the last.fm recommendation radio, which are not necessarily liked by the user. We computed the eclecticism level setting k as 5% and 10%; the results are presented in tables 2 and 3.

Table 2. Users' eclecticism level (k=5)

Number of clusters played	Users' percentage
1	15.28%
2	22.46%
3	28.41%
4	27.69%
5	6.15%

Table 3. Users' eclecticism level (k=10)

Number of clusters played	Users' percentage
1	26.56%
2	34.25%
3	27.08%
4	11.38%
5	0.72%

The results show that the 15.28% (26.56% with k=10) of the users seem to be express a very low level of eclecticism; we can also notice that the percentage of very eclectic users (number of clusters played = 5) is very low: 6.15% (0.72% with k=10). It is important to highlight that this result can be considered as a eclecticism measurement relative to the genres and popularity threshold chosen: the users could actually have played songs belonging to other genres (respect to the five considered) and/or songs having a low level of popularity and consequently not selected (see 2.3 for further details). We also found that, as we are going to explain in details, the eclecticism level is not uniformly distributed among all the category of users.

Who are the less and the most eclectic users? To answer this question, we computed, for each user, the cluster most played, determining five different category of users (classic rock lovers, jazz lovers, pop lovers and so on); then, for each user category, we calculated how the belonging users are distributed in terms of eclecticism level. The results are presented in table 4, for k=5.

Table 4. Users' eclecticism level according to genres loved

Genre loved	Percentage of users per eclecticism level				
	1	2	3	4	5
Metal	32.98%	25.77%	18.56%	18.04%	4.64%
Classic rock	6.82%	20.45%	31.82%	31.82%	9.09%
Pop	9.09%	25.30%	29.25%	30.04%	6.32%
Jazz	4.65%	11.63%	32.56%	39.53%	11.63%

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WebSci '11, June 14-17, 2011, Koblenz, Germany.
Copyright held by the authors.

Hip-hop	14.93%	22.82%	31.27%	26.48%	4.51%
---------	--------	--------	--------	--------	-------

The results show that, at least for the data considered, the Heavy Metal lovers are the less eclectic category of users (32.98% of them has eclecticism level 1, more than twice respect to the global percentage, 15.28%, showed in table 2), while the Jazz lovers seems to be the most eclectic ones, which percentage, for levels 3, 4 and 5, always over the ones showed by the other categories of users.

5. CONCLUSIONS AND FUTURE WORK

Using a the listening history of a sample set of last.fm users and the relative music content description provided by the last.fm folksonomy features, we analyzed the relationship between a folksonomy-based music classification and a collaborative filtering based music classification. We found that, for five music genres selected, there is an high correlation between the song clusters computed using the former approach and the ones computed using the latter.

Clustering the songs also allowed us to compute the eclecticism level of users - in term of diverse collaborative-filtering clusters frequently visited by users - and to find that some categories of users seems to be more prone to eclecticism respect to others e.g. Jazz lovers seems to be much more eclectic than Heavy Metal lovers.

We know that the results found could be dependent on the music genres selected; for this reason we have planned to further investigate this phenomenon, taking into account all the most

popular genres represented in last.fm. This future extension of the study can also work as a base to develop a global collaborative filtering map of the music.

6. REFERENCES

- [1] Szomszor, M., Cattuto, C., Alani, H., O'Hara, K., Baldassarri, A., Loreto, V. and Servedio, V. D. P. (2007) Folksonomies, the Semantic Web, and Movie Recommendation. In: 4th European Semantic Web Conference, Bridging the Gap between Semantic Web and Web 2.0, 3-7th, June 2007, Innsbruck, Austria.
- [2] Yi Zhen, Wu-Jun Li, and Dit-Yan Yeung. 2009. TagiCoFi: tag informed collaborative filtering. In Proceedings of the third ACM conference on Recommender systems (RecSys '09). ACM, New York, NY, USA, 69-76.
- [3] M. Sordo, O. Celma, M. Blech, and E. Guaus. The quest for musical genres: Do the experts and the wisdom of crowds agree? In Proc. of ISMIR, Philadelphia, PA, USA, pp. 255-260, 2008.
- [4] G. Geleijnse, M. Schedl and P. Knees. The quest for ground truth in musical artist tagging in the social web era, in Proc. of the 8th International Conference on Music Information Retrieval, pp. 525-530, 2007.
- [5] D. Ellis, B. Whitman, A. Berenzweig, and S. Lawrence, "The quest for ground truth in musical artist similarity," in Proc. of the International Conf. on Music Information Retrieval, 2002.