# Music Discovery with Social Networks

Cédric S. Mesnage
Faculty of Informatics
University of Lugano
Lugano, Switzerland
cedric.mesnage@usi.ch

Asma Rafiq
Centre for Digital Music
Queen Mary University of London
London, UK E1 4NS
a.rafiq@qmul.ac.uk

Simon Dixon
Centre for Digital Music
Queen Mary University of London
London, UK E1 4NS
simon.dixon@eecs.qmul.ac.uk

Romain P. Brixtel
University of Caen
France
rbrixtel@info.unicaen.fr

## ABSTRACT

Current music recommender systems rely on techniques like collaborative filtering on user-provided information in order to generate relevant recommendations based upon users' music collections or listening habits. In this paper, we examine whether better recommendations can be obtained by taking into account the music preferences of the user's social contacts. We assume that music is naturally diffused through the social network of its listeners, and that we can propagate automatic recommendations in the same way through the network. In order to test this statement, we developed a music recommender application called Starnet on a Social Networking Service. It generated recommendations based either on positive ratings of friends (*social recommendations*), positive ratings of others in the network (*non-social recommendations*), or not based on ratings (*random recommendations*). The user responses to each type of recommendation indicate that social recommendations are better than non-social recommendations, which are in turn better than random recommendations. Likewise, the discovery of novel and relevant music is more likely via social recommendations than non-social. Social shuffle recommendations enable people to discover music through a serendipitous process powered by human relationships and tastes, exploiting the user's social network to share cultural experiences.

## Categories and Subject Descriptors

H.5.5 [**Information Systems**]: Information Interfaces and Presentation—*Sound and Music Computing*; H.4.3 [**Information Systems Applications**]: Communications Applications—*Internet*; H.3.3 [**Information Systems**]: Information Storage and Retrieval—*Information Filtering, Selection Process*; H.3.4 [**Social Networking**]:

*WOMRAD 2011* 2nd Workshop on Music Recommendation and Discovery, colocated with ACM RecSys 2011 (Chicago, US)


## General Terms

Design, Experimentation, Performance

## Keywords

Music discovery, music recommender systems, social networking services

## 1. INTRODUCTION

An interesting mechanism of music discovery occurs when family and friends recommend each other music that they discover. The emergence of Social Networking Services (SNS) and Web Music Communities (WMC) provide us the opportunity to develop music recommender applications to support this mechanism. SNS are becoming increasingly popular means for people to socialise online. Music is playing a similar role on these platforms as in real life social networks. It is shared, discussed, recommended and discovered with social contacts. It is noteworthy that music information seeking behaviour has been indicated as 'highly social' [5]. This social aspect of music should be incorporated in current music recommender systems. WMCs like Last.fm[1], Pandora[2] and Ping[3] are playing a vital role in helping music listeners to build relationships with similar music-listeners and get recommendations based on their current music collections. WMCs are very popular among music fans.

Music discoveries often result from passive behaviour [5]. In [2] authors also indicated that music discovery was seldom a conscious activity until the project participants were given the task of writing diaries when they encountered new music. Therefore, it is quite likely that many people are interested in discovering music but not actively seeking for it on WMCs. Possibly due to the region-specific content access restrictions, since sites such as Pandora can only be used within U.S. territories and Last.fm radio is not available without paid subscription to countries other than UK, US and Germany. On the other hand, SNS such as Facebook[4] allow social interaction with family and friends around the globe.

In this work, we model that music discoveries take place

[1] `http://www.last.fm`
[2] `http://www.pandora.com`
[3] `http://www.apple.com/itunes/ping`
[4] `http://www.facebook.com`

via natural diffusion of music through social networks or randomly. An experiment was conducted to reproduce this process so that we could analyse how people respond to the recommendations. The recommendations arising from such processes are either randomly picked from the pool of tracks of the data set or collaboratively from the tracks recommended by other people on the SNS. A successful music discovery occurs when the user of the application likes a track that s/he has never heard before.
This paper has been divided into 7 sections; section 1 is the introduction, section 2 elaborates the rationale behind this research experiment. In section 3, we discuss the methodology. In section 4, the results are presented. In sections 5, 6 and 7, limitations are discussed, the research work is concluded, and future work is proposed respectively.

## 2. BACKGROUND AND MOTIVATION

Research on finding new music shows that music discovery often occurs with the personal acquaintances playing music to the respondents, and that social networks continue to play a vital role in music discovery in the digital age [10]. Our social contacts may influence our music preferences [3]. Social context plays a significant role in improving music recommendation algorithms [7]. Music can both reflect and define social identity and membership in a given subculture [5]. Information retrieval from social media aids the collection, storage and review of music of the users. It presents opportunities for improved music recommender systems incorporating music preferences of the user's social contacts. An online survey conducted by Entertainment Media Research Company (EMRC) and Wiggin (2009) with 1,608 participants from the United Kingdom, indicated that social networking sites are frequently used for music streaming [1]. Although Last.fm uses collaborative recommendation algorithms, it does not explicitly provide an option to restrict recommendations to the user's social contacts rather than the whole WMC. Interestingly, Pandora has recently attempted to add the feature of "Music Feed" on Pandora One, which shows the activities of friends such as likes, tracks they are listening to and comments. It is a similar concept to Ping but is only available to the paid subscribers and is currently in testing phase [9]. The number of active users on Facebook[5] outnumbers any of the WMCs mentioned above by a significant margin. Therefore, it is more likely to find real life friends on Facebook as compared to the WMCs, forming another motivation to conduct our experiment on Facebook. In [4] authors show that collaborative filtering based on social relationships and tags outperforms standard information retrieval techniques by running simulations on users' listening history. Other research has shown that music listeners sometimes enjoy randomly ordered recommendations [6]. We approach the problem with a different methodology by conducting a live experiment in which users listen and rate the recommendations.

## 3. METHODOLOGY

The aim of the experiment is to test that discoveries are diffused through the social network. The problem is treated as a recommendation problem defined as follows: given a pool of items, select an item that the subject has not heard and that is relevant to her/him. A collaborative recommendation makes use of items rated previously by other subjects to choose the item to be recommended. If collaborative recommendations based on ratings by people from the social network of the subject lead to more successful recommendations than ratings from people not in the subject's social network, then there is an indication that social recommendations are more appropriate for collaborative recommendations.

### 3.1 Experimental setting

The idea of the social shuffle is to recommend tracks (see Section 3.2 for more details) and diffuse discoveries through the social network. Figure 1 shows an example of this process when a recommendation is posted by a user in her/his SNS. In this case, Joe gets a random recommendation, gives it a 4 star rating, the track is then diffused to his social network. If a friend of Joe gives a high rating to the same track, it will be diffused to her network as well (in this case, the example of Alice). If Joe's friend does not enjoy the track and gives it 2 or less rating then the social diffusion stops and here Aleks' friends will not get this recommendation.

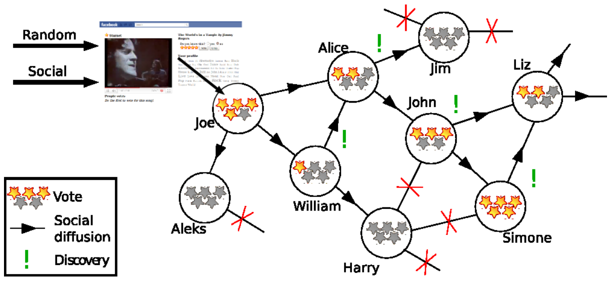


**Figure 1: Example of the social shuffle principle.**

Each time a track is recommended, the subject of the experiment rates the track on a 0 to 5 Likert scale; a psychometric scale introduced by Rensis Likert [11] visually represented as stars, where selection of 0 indicates dislike and 5 indicates favourite tracks.

#### 3.1.1 Dataset

In order to have a representative pool of available tracks for random selection, the dataset was built from Last.fm containing a million tracks fetched through their API. We explored a subset of the Last.fm of about 300,000 people and their friends. The tag profile of each user was fetched, giving us the list of tags they used to organize tracks, yielding a set of more than 300,000 unique tags. For each tag the API allows to fetch the top 50 most popular tracks tagged with this tag. Only the set of tracks is relevant for this study.
On Facebook, we started by fetching the network of the first author, his friends and friends of friends interested in participating in this experiment. Many people share music using Youtube[6] by posting links on Facebook. The advantage of using music videos from Youtube is that the full track can be played (if the video is available), which is not the case for public users of Last.fm for instance which limits the playback to 30 seconds. The limitation pertaining to Youtube videos is that sometimes these are blocked in various countries and

[5]http://www.facebook.com/press/info.php?statistics

[6]http://www.youtube.com

removed for copyright violations. The application detects when such errors occur and does not recommend that video afterwards. On the Youtube API, videos were searched for each track with the artist name and track title. We selected the most popular video for each track and restricted to fetch the videos tagged as “Music” only. Using this process about a quarter of the tracks from Last.fm had a video on Youtube, totaling around 252,000 tracks.

### 3.1.2 *Facebook application*

Starnet is a Facebook application fed by ratings on random selections. A positive rating spawns diffusion through the user’s social network, i.e. the shuffle recommendation becomes social. The interface (Figure 2) consists of the current track description (title and artist name), the music video associated, a tag cloud of the user’s profile and a rating form consisting of 5 stars. Also, it has “next” button to play the next track and a “bail” button which sets the stars to 0 and plays the next track. The user is asked to indicate whether the track is already known, in order to learn whether it is a music discovery. We assume it is a reasonable estimate for true unknown tracks.

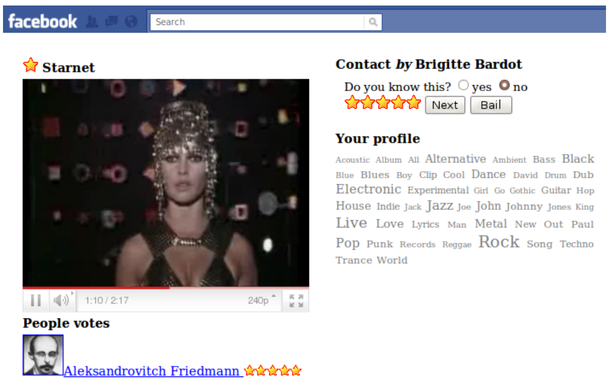


**Figure 2: Screenshot of Starnet Application on Facebook.**

### 3.1.3 *Subjects*

The subjects are the people who used the Starnet application. A total of 68 subjects participated in this experiment and allowed access to their social profiles. Each subject acted as his/her own control group by getting and rating random recommendations. In about 4 months (from the 29th June 2010 to the 18th October 2010), 31 subjects made 4966 ratings using the Starnet Application. Participation of other 37 subjects was insignificant in terms of ratings.

## 3.2 Recommendation Strategy

The recommender system produced the following types of recommendations:

- Random recommendations: The random selection is a query which selects tracks that have not been rated by the subject and orders them randomly.

- Collaborative recommendations: Selects a track randomly from the set of tracks that have been rated with a rating above average (i.e. greater than 2 stars) and not yet rated by the subject. The collaborative recommendation can be social or non-social.

  - Social recommendations: The social recommender selects a track randomly, from the tracks that have been rated by friends of the subject, with a rating superior to 2 stars.

  - Non-social recommendations: The non-social recommender selects a track randomly, from the tracks that have been rated by people who are not friends of the subject, with a rating superior to 2 stars.

The likelihood for selecting random recommendation, social recommendation and non-social recommendation is 0.5, 0.25 and 0.25, respectively.

# 4. RESULTS

This section presents the results from the analysis of the ratings made by the subjects of the experiment. The success of a recommendation model to discover new and relevant tracks for a subject can be evaluated without further user input.

*Discoveries.*

A measurement of the success of a recommendation model to discover new and relevant tracks for a subject is the ratio of already discovered tracks and all rated tracks. Figure 3 represents how these sets relate for a particular user.

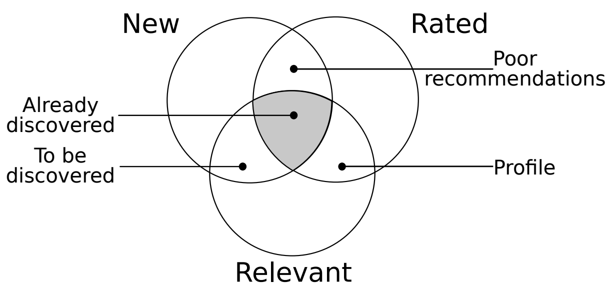


**Figure 3: Discovery recall diagram.**

Histograms are used to compare distributions of ratings. The $x$ axis represents number of stars/rating and the corresponding percentage of ratings on the $y$ axis. Heat maps are used to show how the ratings of a subject relates to ratings on the same tracks from friends and non-friends. The ratings for social and non-social recommendations of known and unknown tracks are compared, showing that social recommendations give better ratings than non-social ones.

Figure 4 shows the ratings for the tracks known to the subject and for those specified as unknown, respectively represented as black and white bars. This shows a clear distinction between known and unknown tracks, most unknown tracks are disliked whereas most known tracks are liked. In this figure we looked indifferently at all the recommenders, in the next figure we differentiate between collaborative and random recommendations.

Figure 5 represents the proportions of ratings for collaborative and random recommendations. Collaborative recommendation outperforms random recommendation as 45% of

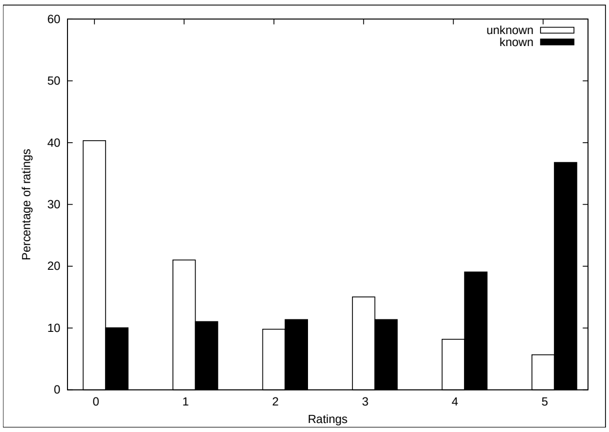


**Figure 4: Histogram of ratings for known and unknown tracks across all the users.**

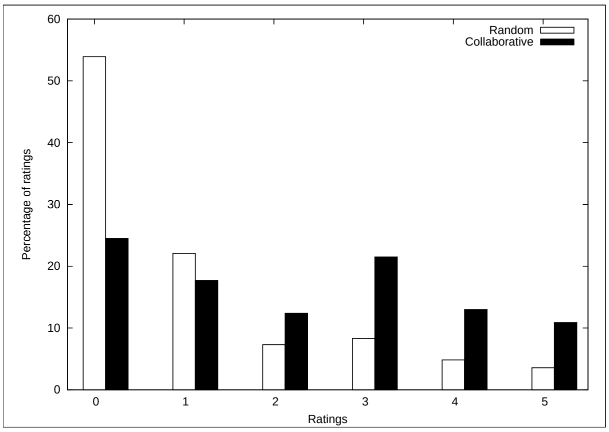


**Figure 5: Histogram of ratings for collaborative and random ratings.**

its recommendation get 3 star or above ratings. The percentage drops to around 17% in the case of random recommendations. In this visualization, an increased number of ratings above the threshold of 3 stars is evident for collaborative recommendations. This is due to the way collaborative recommendations are made: only tracks which get a rating superior to 2 stars from other users are used as collaborative recommendations. This shows that there is some consistency in users' ratings. We now look only at the collaborative ratings dividing them between social and non-social recommendations.

Figure 6 represents the ratings from two types of collaborative recommendations. The social and non-social ratings have different distributions indicating that subjects react differently to recommendations coming from their friends than from people they do not know. The subjects are not aware of the source of the recommendation. More than 47% of social recommendations get 3 or more stars as compared to the non-social ones, (33%). In this visualization, we see that social recommendations tend to lead to better ratings, but we mix known and unknown tracks.

Figure 7 represents the ratings of collaborative recommendations on known and unknown tracks. 94% of the ratings state that the track is unknown to the subject, which is why the left histogram on Figure 7 is similar to the histogram of Figure 6. In general, the social recommendations lead to more unknown good recommendations (46% at least 3 stars against 34%) than the non-social ones and less bad recommendations (40% at most 2 stars against 60%). The histogram on the right shows ratings where the subject specified she knew the track. Although, the overall rating distributions are very different, the same trend of higher ratings for social recommendations is evident.

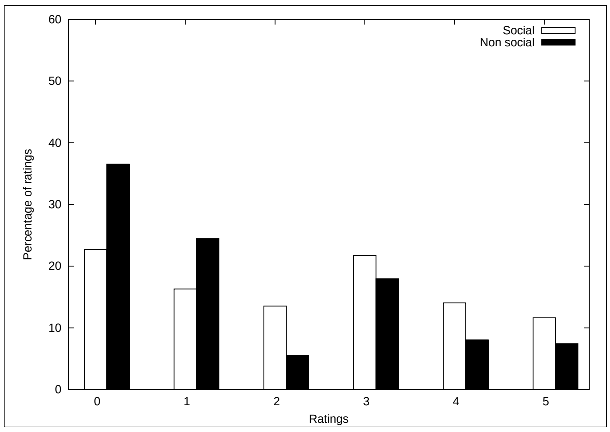


**Figure 6: Histogram of social and non-social ratings.**

Figure 8 shows two heat maps. Each square represents the proportion of ratings made by all users of the application on the same tracks for each rating. The left hand heat map represents the relation between ratings of people who are not friends and the right one the relation between ratings of people who are friends.The contrasted regions on both heat map shows that people agree on ratings on the same tracks. Interestingly, the social heat map is more contrasted, showing that friends agree more with their ratings than people who are not friends.

## 5. LIMITATIONS

Our results show that people tend to prefer the music that their friends prefer. One of the limitations of this work is that it is based on the assumption that the relationships in SNS or WMC are with real life friends and family which might not be true for some subjects. Therefore, they might not be very good source for music recommendation or discovery in that case, as they might not share user's music taste. Results suggest otherwise.

Studies on music behaviour also suggest that sometimes youngsters try to distance themselves from the previous generation [8] in which case, social recommendations from older members of the social contacts might not work well. So, there should be an option for tuning the users' social network for music recommendation, when such a feature is integrated in a music recommender system.

## 6. CONCLUSION

The analysis of the results of the experiment on social diffusion. lead us to the following conclusions:

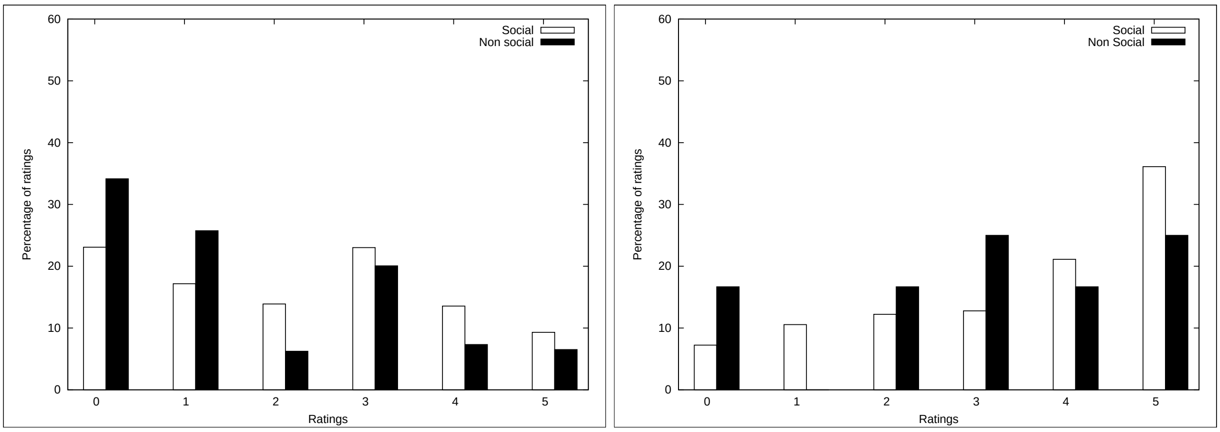


**Figure 7: Histograms of ratings of social and non social recommendations on tracks unknown (left) and known (right).**

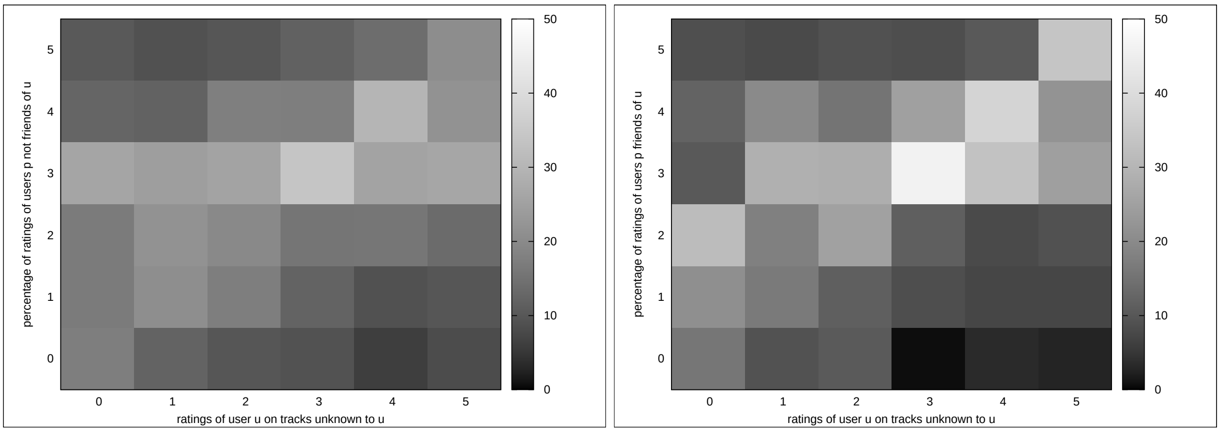


**Figure 8: Heat map of ratings made on the same tracks, by people who are not friends and did not know the track on the left and people who are friends on the right.**

- Social recommendations were preferred by the users over the non-social recommendations which indicates that people tend to share music taste with their friends on Social Network Services than with people who are not their friends.

- Recommended tracks that were known are mostly liked whereas recommended tracks that are unknown are mostly disliked.

- Collaborative recommendations lead to better recommendations than random recommendations.

These conclusions support the view that social diffusion is a good mechanism for music recommendation and discovery. It is anticipated that it shall form the foundation for the framework of better music recommender systems combined with social media.

## 7. FUTURE WORK

An improved recommendation system could make use of the genres (or tags) of the tracks previously rated by the subject. We have been working on a different input for music videos. We are now fetching the Youtube videos posted on Facebook by registered people to the Starnet application and their friends. This changes the settings of the application and requires us to focus on interaction mechanisms for people to explore a social network. The new application shall allow users to add and remove social contacts to their network, to tune their personalised social radio.
Another application of interest is a music recommender system that recommends music based on the type of event and music preferences of the people attending it. The initial prototype of the system is available as a Facebook Application[7]. People post events (such as party, wedding, etc.) on SNS. Analysis of music taste of the attendees of the events can enable the event organiser to make better decisions on what type of music shall be enjoyable for most of the attendees. However, event categorisation and determination of suitability of music within that particular category can be a challenge.

[7] `http://apps.facebook.com/music_valley/`

# 8. REFERENCES


[1] Z. Baker. *Social Networking and Music: the impact of myspace as a Digital Platform for Discovery and Distribution.* Bachelor's thesis, Cardiff School of Journalism, Media and Cultural Studies, Cardiff University, 2009.

[2] S. Cunningham, D. Bainbridge, and D. McKay. Finding new music: A diary study of everyday encounters with novel songs. pages 83–88. Austrian Computer Society, International Society for Music Information Retrieval, 2007.

[3] T. DeNora. *Music in Everyday Life*, pages 4–8. Cambridge University Press, 2000.

[4] I. Konstas, V. Stathopoulos, and J. Jose. On social networks and collaborative recommendation. In *SIGIR '09: Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 195–202, New York, NY, USA, 2009. ACM.

[5] A. Laplante. *Everyday life music information-seeking behaviour of young adults: an exploratory study.* PhD thesis, McGill University, Montreal, Canada, 2008. pages 4-65.

[6] T. Leong, F. Vetere, and S. Howard. The serendipity shuffle. In *OZCHI '05*, pages 1–4. Computer-Human Interaction Special Interest Group, 2005.

[7] D. McEnnis and S. Cunningham. Sociology and music recommendation systems. pages 185–186. Austrian Computer Society, 8th International Conference on Music Information Retrieval, 2007.

[8] A. Minks. Growing and grooving to a steady beat: Pop music in fifth-graders' social lives. *Yearbook for traditional music*, 31:77–101, 1999.

[9] M. Siegler. Dancing to the drumbeat of HTML5, New Pandora is brilliant, beautiful. Website, 2011. `http://techcrunch.com/2011/07/12/new-pandora/`.

[10] S. Tepper and E. Hargittai. Pathways to music exploration in a digital age. *Poetics*, 37(3):227–249, 2009.

[11] J. Uebersax. Likert scales: dispelling the confusion. Website, 2006. `http://john-uebersax.com/stat/likert.htm`.