Meerkat: Exploring Semantic Music Discovery Using Personalized Radio

Ashley M. Oudenne
Swarthmore College
Swarthmore, PA
aoudenn1@swarthmore.edu

Youngmoo E. Kim
Drexel University
Philadelphia, PA
ykim@drexel.edu

Douglas S. Turnbull
Swarthmore College
Swarthmore, PA
turnbull@cs.swarthmore.edu

ABSTRACT

We present a Internet radio player called *Meerkat* to explore the use of semantic tags as a mechanism for controlling a personalized stream of music. We argue that this increases the level of interactivity over standard commercial Internet radio players and, as a result, produces a positive and efficient music discovery experience. We also challenge the common design decision not to reveal the identity of upcoming songs. We argue that showing these songs to the user provides context for the current state of the radio station. More importantly, it gives the user instant visual feedback when he or she decides to dynamically alter the radio station. Based on a small initial user study, we find evidence to suggest that showing the upcoming songs improves the user experience.

Keywords

Semantic Music Discovery Engine, Personalized Internet Radio

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: Evaluation, User-Centered Design; H.5.5 [**Sound and Music Computing**]: System

1. INTRODUCTION

In the last decade, there has been an enormous commercial and academic interest in developing web and mobile *music discovery engines* [1]. These engines help billions of music fans discover music from millions of artists. While there are many design paradigms for music discovery, one of the popular paradigms is the personalized Internet radio model (e.g., Pandora, Slacker). In general, a user will start a *music station* by selecting a favorite artist. The user will then *control* the station by indicating that they approve or disapprove of the current song. While the level of interactivity is somewhat limited, it is the simplicity of this interactive design that has made personalized Internet radio popular.

In this paper, we explore the use of semantic tags for the purpose of controlling a personalized radio station. That is, users select tags, such as "bluegrass," "raspy female vocals," and "energetic," to indicate the type of music they wish to hear. Our system, called *Meerkat*[1], dynamically adapts the playlist to reflect the selected tags. It also displays tags that describe the current song. These song tags both provide context for the current state of the music station and suggest novel tags that the user may want to use to control the station.

While our primary goal is to explore the use of semantic tags as a control mechanism for music discovery, we are also interested in testing a standard design decision of traditional personalized Internet radio interfaces: hiding the identity of upcoming songs. In a commercial setting, this design decision is directly related to licensing contracts that limit the freedom of the user to have "on-demand" access to music. However, the growing supply of copyright-cleared or copyright-free music available online may decouple us from this commercial limitation in the future, and the decision to either show or hide the upcoming songs may become an important decision when designing a music discovery interface.

2. MUSIC DISCOVERY ENGINES

Music discovery engines, which allow users to discover new songs, artists, or types of music, can be generally categorized into three paradigms. The first is the *encyclopedia* paradigm, in which a user browses music using relational data (e.g., Last.fm[2], AMG All Music[3]). The second paradigm is the *metadata search*, in which a user queries a database using a specific song, artist, or album name (e.g., Apple iTunes Store[4], Deezer[5]). The third paradigm, the *radio* model, is one in which a user serendipitously discovers music while listening to a stream of music (e.g., Pandora[6], Slacker[7]). We note that these three paradigms are not mutually exclusive. For example, Last.fm incorporates elements of all three in its user interface.

The radio model is particularly compelling for music discovery because it is passive and unpredictable. That is, a listener can simply listen to the radio and serendipitously discovery new music. The model is derived from the traditional AM/FM radio experience, in which a DJ constructs a playlist of songs that are revealed over time. Often, the DJ will release hints about upcoming songs both to give context for the station and to entice the listener to keep listening. This model was quickly extended to streaming radio via the Internet so that listeners could have access to prepared playlists

[1] http://mugwort.cs.swarthmore.edu/smir/pages/meerkatDemo.html
[2] http://www.last.fm
[3] http://www.allmusic.com
[4] http://www.apple.com/itunes
[5] http://www.deezer.com
[6] http://www.pandora.com
[7] http://www.slacker.com

from stations from around the world using Internet radio portals (e.g., SHOUTcast[8], AccuRadio[9]).

More recently, there has been a growing interest in *personalized* Internet radio. It is similar to the traditional radio paradigm in that the user listens to a stream of music, but different in that the user is able to *control* his or her own radio station using interactive design features. For example, using Pandora, a user can start a radio station by selecting a favorite artist. Based on this seed artist, Pandora will begin streaming music by similar artists. The user will see metadata (song name, artist name, album cover artwork) associated with the current song as well as the two previous songs, but will not see any information associated with future songs. The user is allowed to fast-forward songs[10] but is neither allowed to rewind the music station nor to replay previous songs. As such, these restrictions limit the music discovery experience. Once a station has been started, a user can indirectly control the radio station by indicating that they like or dislike a song. However, while this is an improvement over traditional radio in terms of interactivity, we argue that this still results in a relatively limited interactive experience.

There are a number of personalized radio player interfaces that provide greater interactive control over a stream of music, but often such interfaces present a more complex (and potentially overwhelming) user experience. For example, Lamere and Eck's "Search Inside the Music" [2] and Lillie's "MusicBox" [3] project songs into a 3D space which can be rotated and rescaled according to metadata and acoustic features. Both interfaces allow a user to create a path through the constructed 3D space. This path effectively represents a music playlist for a personalized radio station. Pampalk and Goto's "MusicSun" [4] allows users to search for music based on three notions of music similarity: acoustic, related webpages, and social tags. There are optional sliders at the bottom of the interfaces that allow users to control the relative weight of importance for the three similarity measures. These "advanced" features keep the interface relatively easy-to-learn while providing extra functionality for power users. MusicSun is the most similar music discovery interface to our Meerkat interface for three reasons: both use a simple wireframe design, both use tags to control the personalized music experience, and both use tags to provide context for the current state of the radio playlist. However, our interface design is more akin to standard internet radio players and is more minimalist because we display less auxiliary information about fewer songs at any given point in time. Finally, we note that there are hundreds of music discovery interfaces, some of which are described on the VisualizingMusic blog[11].

3. MEERKAT

We are interested in gaining a better understanding of music discovery so that we might improve the user experience in the future. In this paper, we specifically explore two questions:

- Are semantic tags useful as a controlling mechanism for personalized Internet radio?
- Does knowing the songs that will be played in the future improve the music discovery experience?

[8]http://www.shoutcast.com/
[9]http://www.accuradio.com/
[10]Due to licensing restrictions, Pandora must limit the number of song that are skipped per hour.
[11]http://visualizingmusic.com

In order to study these (and related) questions, we have created an online music discovery system called Meerkat (see Figure 1). The back-end uses four sources of music information (web pages, social tags, audio content, and preference information) to index music with a large set of semantic tags [5, 6]. Our current music corpus consists of 10,870 songs that are representative of 18 genres (e.g., "rock". "electronic', "classical") and 180 subgenres (e.g., "grunge", "trance", "romantic period opera"). Our vocabulary of tags consists of hundreds of genres & subgenres, dozens of emotions ("happy"), dozens of instruments ("distorted electric guitar"), hundreds of acoustic characteristics ("dominant bass riff"), and thousands of free-text social tags (from Last.fm). Based on a list of tags, the back-end returns a ranked list (i.e., a playlist) of semantically-relevant songs.

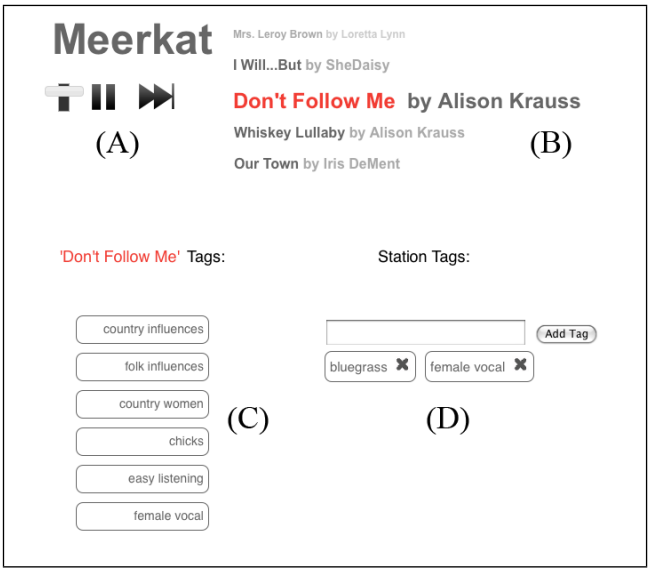

Figure 1: Meerkat Screenshot: (A) The volume control, play/pause toggle, and fast-forward button allow the user to control the music playback. (B) Past songs (in a gray font) are displayed above the current song (in a red font) while future songs (also in a gray font) are below. (C) A list of tags associated with the current song provides context for the music. A user can click on a tag if he or she wishes to add it to the radio station. (D) The user selects tags to control the radio station. A user can also add tags using the text box (with automatic tab completion).

The Meerkat front-end follows the personalized radio paradigm: a playlist of relevant songs is played one at a time based on input from the user (see Figure 1). Meerkat is somewhat different from most music discovery engines in that it allows users to explicitly control the radio using multiple tags. For example, a user might start a radio station by selecting "instrumental," "bluegrass," and "accordion." The user might then remove the tag for "accordion" and replace it with "banjo" to dynamically alter the stream of music as he or she sees fit. When a tag is added or deleted, the upcoming songs change to reflect the newly generated radio playlist. Meerkat's user interface was deliberately kept as a wireframe so that the functionality of the system, and not its layout or design, was emphasized to users. In order to isolate the effects of "seeing the future" in a music discovery interface, the upcoming songs could be displayed or hidden from users. If the upcoming songs were hidden and a user added or deleted a tag, he or she received

no visual feedback about the change in the type of music the radio would play next. This "No Future" mode is the standard design paradigm for most personalized Internet radio players.

4. EXPERIMENTAL DESIGN

We designed a one-week user study in order to test the interface. Fifteen participants were asked to use the interface for about an hour a day whenever they would normally listen to music. We split these participants into two groups: those who could see the future, and those that could not. There were seven participants in the "Future" group and eight participants in the "No Future" group.

Our user study was conducted using mostly college-age volunteers with varying degrees of computer proficiency. Of the user study participants, seven identified as male, seven as female, and one as "other". All participants felt very comfortable using the Internet and all but one had used some sort of online music player such as Pandora Internet Radio or radio streaming via the Internet. To make the experience as natural as possible, using the interface was not enforced, so some participants listened more than the suggested time, and others listened less.

During the user study, we collected two types of data. The first, the *explicit* data, was collected by survey at the beginning and end of the user study. In the surveys, we attempted to determine testers' familiarity with online music players, the features they felt were most important to a music discovery user interface, and their experiences while using Meerkat. The second, the *implicit* data, was collected by tracking each click the tester made while using Meerkat. We also tracked the length of each user's Meerkat session, the time spent listening to each song, and the number of songs the user finished or skipped through. By combining these two types of data together, we were able to extract meaningful information about both the experience of using tags to control the radio and the effects of revealing the identity of future songs on music discovery.

We recorded two important implicit statistics during the Meerkat user study: the number of songs skipped and the number of songs finished. We define a skipped song to be any song not finished by the user, except for the song to which the user was listening when he or she ended a Meerkat session. We also documented the number of songs skipped before and after ten seconds, reasoning that ten seconds gives users enough time to decide whether or not they enjoy listening to the current song. If they skip before the ten seconds are finished, they may be searching for particular songs rather than attempting to discover new music.

Another important statistic is the number of songs finished. For convenience, we count the song to which users were listening when they ended a Meerkat session as being a finished song. Finished songs indicate that the user was likely listening to a song that he or she enjoyed.

5. RESULTS AND DISCUSSION

The majority of the participants used Meerkat for five to eight hours during the weeklong user study and discovered between one and four new artists and between five and ten new songs that they liked.

5.1 Explicit Results and Discussion

Overall, the test subject response to the interface and the future songs feature was very positive. Users enjoyed using Meerkat and particularly liked the easy-to-use interface. Our data indicates that testers enjoyed using semantic tags to control the radio. It also suggests that showing the future songs is not only a feature that users want to have in a music interface, but is also one that helps them discover new music.

Table 1: Using Tags to Control Radio

	Liked	Didn't Care	Disliked
No Future Group	7	0	1
Future Group	7	0	0

Testers responded positively to using tags to control their radio stations. In both the Future and No Future groups, all but one tester liked using tags (see Table 1). This suggests that, for a music discovery interface, tags may be a good way of controlling the radio station.

Table 2: Seeing the Future

	Liked	Didn't Care	Disliked
No Future Group	4	3	1
Future Group	7	0	0

Tester response to seeing the future songs was varied (see Table 2). The No Future group felt lukewarm toward seeing the upcoming songs. Although they did not have this feature enabled during their Meerkat sessions, they were given a brief chance to try it while they took the concluding survey. These users did not believe that seeing the upcoming songs would greatly improve their experience. However, testers in the Future group all agreed that they liked seeing the future songs. This suggests that users may not realize how much they enjoy seeing the future songs until they get a chance to interact with them for a prolonged period of time.

Table 3: Amount of New Music Discovered with Future Enabled

	More	Same	Less
No Future Group	2	3	3
Future Group	5	1	1

We also asked users about the amount of music that they discovered or would have discovered if the future songs were enabled (see Table 3). Testers in the No Future group again had varied responses, some feeling that they would have discovered more music than they did, some less, and some the same amount. Most of the Future group felt that they discovered more music with the future songs enabled. This suggests that users who have access to this feature generally feel that it aids in the music discovery process.

Additionally, the majority of users in the Future group found their overall experience while using Meerkat to be very enjoyable and fun, which indicates a positive user experience (see Table 4). Subjects in the No Future group had mixed responses ranging from fun to boring.

Table 4: Overall Experience Rating

	Had a Great Time	OK	Boring
No Future Group	1	5	1
Future Group	6	2	0

Based on our small study, hiding the future songs actually has a noticable impact on the user experience, possibly because the user receives no system feedback when they add or remove a tag from the radio tag list. This in turn decreases the transparency of the system, making it less interactive. Based on our results, it seems that showing the future is not only desirable to users but also makes the interface more fun to use. This observation is contrary to the common design practice of not showing the title of the upcoming songs.

5.2 Implicit Results and Discussion

The fifteen test subjects used Meerkat for a combined total of seventy sessions and listened to 2,527 songs. They completed 1,285 of these songs. The seven participants in the Future group listened to 1,225 songs and the eight participants in the No Future group listened to 1,302 songs. We averaged the following statistics using the total number of songs played.

Table 5: Tags Added Per Song

	Tags Added
No Future Group	0.2204
Future Group	0.2688

The first statistic that we extracted from user sessions is the tags added per song. According to the data, the No Future group added 0.0484 fewer tags per song than did users who could see the future (see Table 5). While this in itself does not indicate a positive or negative response to using tags to control the radio, it does indicate that users in the Future group were more likely to add tags than users in the No Future group. This may be caused by the fact that the Future group received immediate feedback from the system when they added a new tag (i.e. the upcoming songs changed) while the No Future group received no visual feedback.

Table 6: Skipped and Finished Song Rates

	Total Skipped Songs	Skipped Before 10 Seconds	Skipped After 10 Seconds	Total Finished Songs
No Future Group	0.5431	0.3078	0.2376	0.4547
Future Group	0.4409	0.2343	0.2066	0.5591

We next considered the skipped and finished songs per song. The No Future group skipped songs 0.1022 more times per song than the Future group (see Table 6). The test subjects without the future also tended to skip 0.0735 more times per song before the end of the ten second trial period. This may indicate that testers without the future were not taking time to try new songs. Since users could skip as many songs as they desired, the anticipation of the unknown next song may have encouraged them to skip the song too early.

The No Future group tended to skip 0.0310 times more often per song after the ten second trial period and also finished 0.1044 fewer times per song than the Future group (see Table 6). This indicates a negative user experience for the No Future group since they neither finished songs nor discovered many songs due to their higher rate of skipping. This is corroborrated by the explicit data: testers without the future had a worse overall experience while using Meerkat. Based on these statistics, we again have evidence to suggest that hiding the identity of upcoming songs may be detrimental to the music discovery experience.

6. CONCLUSION

The radio paradigm of serendipitous music discovery is a compelling model. Its allure lies in the listener's inability to predict the future, creating a tension that encourages him or her to continue listening to discover the next song. However, obscuring the future, which makes the model so captivating, may not in the best interest of the listener. When the future is hidden, users may be more prone to skip songs without taking an adequate amount of time to sample the song, and they are less likely to finish a new song. Showing the future makes the interface more fun to use from the user perspective, and the user has a better overall music discovery experience if the future songs are displayed.

Additionally, using tags to control the radio station is a feature that users enjoy because it gives them increased control over the music that they hear. Since the supply of copyright-free or -cleared music is growing, it is now possible to couple semantic tag control and future song displays into a powerful interface that gives users the control that they desire.

7. REFERENCES

[1] P. Lamere and O. Celma. Music recommendation tutorial. International Conference on Music Information Retrieval (ISMIR), September 2007.

[2] P. Lamere and D. Eck. Using 3D visualizations to explore and discover music. *International Conference on Music Information Retrieval (ISMIR)*, 2007.

[3] A. Lillie. Musicbox: Navigating the space of your music. Master's thesis, M.I.T., 2008.

[4] E. Pampalk and M. Goto. Musicsun: A new approach to artist recommendation. In *International Conference on Music Information Retrieval (ISMIR)*, 2007.

[5] B. Tomasik, J. Kim, M. Ladlow, M. Augat, D. Tingle, R. Wicentowski, and D. Turnbull. Using regression to combine data sources for semantic music discovery. *ISMIR*, 2009.

[6] D. Turnbull. *Design and Development of a Semantic Music Discovery Engine*. PhD thesis, UC San Diego, 2008.