Exploring Control and Feedback for Personalized Internet Radio

Richard Roberts, Kristofer Stensland, Douglas Turnbull
Department of Computer Science
Ithaca College
{rrobert1,kstensl1,dturnbull}@ithaca.edu

ABSTRACT
Personalized Internet radio players like Pandora and Last.fm provide users with customized streams of music. While they are extremely popular, these systems are limited in the number of the ways that a user can control the radio station. They tend to provide the user with only a small amount of feedback after the user alters the station. We introduce a system called *MegsRadio* that provides a user with a broad variety of control and feedback mechanisms. Based on a small user study, we find that users make use of many of the novel control mechanism and develop a better understanding of the playlist algorithm based on the feedback they receive from the system.

Keywords
Personalized Internet Radio, Music Discovery, Human-Computer Interaction, User Interface Design

INTRODUCTION
Personalized Internet radio allows a user to *control* a customized stream of music. That is, unlike terrestrial radio (AM/FM stations) or streaming Internet radio (e.g., AccuRadio) which are simultaneously broadcast to a large group of users, personalized Internet radio players create individualized streams of music for each user. Commercial players, such as those created by Pandora and Last.fm, are extremely popular and have millions of daily users. They are often easy-to-use, available for free, and play songs from a large catalog of music.

Personalized Internet radio players are different from *celestial jukeboxes* [1] like Apple iTunes, Rhapsody, and Spotify, in that they do not allow a user to directly select individual songs to be played in a specific order. Rather, a playlists of songs is generated in real-time by a computer algorithm based on input from the user. The most common forms of input are a seed artist or song to start a station and preference ratings (e.g., Thumbs Up / Thumbs Down) for songs that are played on the station. While this allows for some level of customization, we argue that there are additional *control mechanisms* that can aid in the music discovery experience. We also suggest that additional visual and auditory *feedback mechanisms* are important for providing transparently and establishing trust in the playlist algorithm.

In this paper, we first examine three popular personalized Internet Radio players by focusing on playlist control and feedback mechanisms. We then introduce a new system called *MegsRadio* that was specifically designed to explore various aspects of Internet radio personalization. This system incorporates common features like seeding the station with an artist but also allows incorporated many additional control mechanisms. For example, the user can seed the station with semantic *tags* that relate to genre, instrumentation, emotions, and other song characteristics. Furthermore, a user can combine tags and artist similarity to create, for instance, a station that plays "*mellow acoustic rock that sounds like The Beatles but not like Oasis*." Other control mechanisms let the user specify song tempo, danceable or energy, or ste2er the overall mix to focus on music by local artists or by obscure artists. Finally, we explore the importance of the feedback presented to the user based on the decisions that are being made by the automatic playlist algorithm.

COMPARING RADIO PLAYERS
In this section, we compare the three most popular personalized Internet Radio players in the United States according to Alexa[1]: Pandora, Last.fm, and Jango. In particular, we focus on how a user can create and adapt a *station* (i.e., a personalized stream of music), as well as the feedback he or she receives when interacting with the station (see Table 1).

The most common feature among these players is that a user can start a station with a seed artist. Some players also let the user start stations with a particular track or genre. Last.fm has a second type of player, which lets users enter one or more (social) tags such as "happy", "distorted electric guitar", and "New York". Once created, the user can control the station indicating whether they like or dislike songs that are played. Both Pandora and Jango let users add additional seed artists after a station has been created. Last.fm, on the other hand, does not let a user add additional tags or artists once at station has been created.

[1] Data recorded Sept 23, 2011 at http://www.alexa.com/topsites/countries;0/US

Music: Cognition, Technology, Society
Cornell University, Ithaca, NY, USA
May 11-13, 2012

Player	Control	Feedback
Pandora	**Seeding:** artist, song or genre stations **Adapting:** can add additional artists or tracks to artist or track stations, song preference rating (like, dislike, ban artist, play more like this)	**Audio:** first song is often by seed artist, replay liked tracks **Visual:** like/dislike icon is highlighted
Last.fm	**Seeding:** artist or multi-artist stations, tag or multi-tag stations **Adapting:** song preference ratings (like, dislike)	**Audio**: none **Visual**: like/dislike icon is highlighted
Jango	**Seeding**: artist stations **Adapting**: add similar artists suggested by player, specify variety in songs / artists, song preference rating (like, dislike)	**Audio**: first song is often by seed artist **Visual**: none
MegsRadio	**Seeding**: artist or tag stations **Adapting**: additional artists and/or tags, song preference ratings (like, dislike, new discovery, confusing recommendation), station focus filters (local music, popular/obscure artists, song repeating), song filters (tempo, danceablility, energy)	**Audio**: first songs is often by seed artist, **Visual**: like/dislike icon is highlighted, upcoming artist collage, local music banner

Table 1. Control and feedback mechanism for three popular personalized Internet Radio Players and our proposed MegsRadio system. "Seeding" refers to input that can be used to start a new station. "Adapting" refers to how a station can be changed after it has been started.

While there are various options for controlling stations, the three players provide only a small about of visual or auditory feedback based on the user's actions. The most common form of feedback is highlighting an icon to indicate the selected preference. Pandora, more so that the other two, almost always selects a song from a new seed artist whenever a new seed artists is added to the station.

There are many other players (Slacker, MeeMix) which incorporate similar control and feedback mechanisms. One notable exception is Musicovery. Musicovery provides users with a two-dimensional *mood* space that opposes dark and Positive and on the X-axis and Energetic and Calm on the Y-axis in a manner that is akin to Russell's Valance-Arousal Space [Russel80]. They also have a *dance* space that opposes tempo on the X-axis and danceability on the Y-axis. In addition, Musicovery lets users select songs from a specific range of years and from a subset of 20 music genres.

MEGSRADIO

MegsRadio is a research-based personalized Internet Radio player that has been designed to explore various aspects of music discovery. In this paper, we focus on the control and feedback mechanisms that have been built into MegsRadio based on feedback from an initial survey of Internet radio users. Our goal is to determine which mechanisms are useful for music discovery and which one might be potentially detrimental to the system in terms of being confusing or adding unnecessary complexity.

MegsRadio features broad variety of control and feedback mechanisms (See Figure 2). The control mechanisms include positive and negative associations to artists or tags, filters for limiting music by tempo, loudness, danceability, and energy, and playlist specifications to focus on local music, popular artists, or obscure artists, and song preference buttons for indicating like, dislike, or a new discovery.

Like most players, we highlight user interface components to reflect changes that the user has made to the station. However, unlike other players, we show users a set of upcoming artists that are updated whenever a change is made to the station. This is intended to provide the user with instant feedback on how his or her choices have affect the playlist algorithm. We also highlight relevant aspects of a song or artist such as whether the artist has an upcoming event in the user's hometown. In addition to the visual feedback, we adapt the playlist so as to place additional emphasis on the users last action. In the case of negative feedback, this means skipping the current song. In the case of positive feedback, this means finding a song that has a strong semantic association to the given action. For example, if the user adds the "mellow" tag, we will find a song that is strongly associated with mellow even if it means down

weighting other tags that the user had previously added to the station.

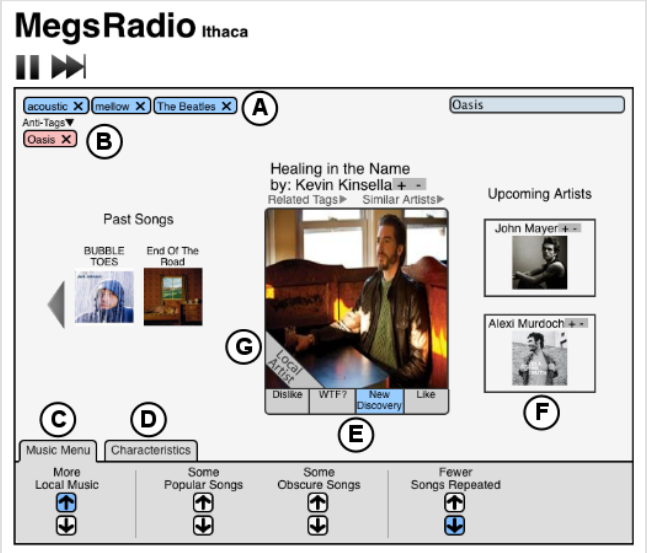

Figure 1. Screenshot of MegsRadio: A user can setup a station by (A) adding artists and tags (e.g., genres, emotions, instruments, etc.) or (B) negative associations with artists and tags. The Music Menu (C) lets users focus the stream of music on popular artists, obscure artists, or local artists, and allows them to select how often they want songs to be repeated. Using Song Characteristics (D), a user can filter tracks by tempo, danceability, loudness, and energy. Users can also indicate (E) whether they like or dislike each song as it plays. Whenever a user adapts the station, the playlist is recalculated and the upcoming artists (F) are updated. Finally, a special banner (G) is displayed if the current song is by a local artist or an artist with upcoming event in the nearby region.

USER STUDY

We conducted a short study to explore the various control and feedback mechanisms described in the previous sections. The study involved 16 undergraduate students at a small liberal artists college in the northeast United States. Student volunteers received extras course credit for their participation. Our initial evaluation consisted of pre-study surveys that focused on music listening habits. In addition, eight students were invited for individual interviews where we discussed their listening habits in more depth and observed them using MegsRadio for the first time.

We then asked the students to spend one week using MegsRadio whenever they would normally listen to music on their computers. During this time, we collected implicit usage data. At the end of the week, we conducted a post-study survey to learn about the overall user experience as well as to gain a better understanding of how like or dislike the various control and feedback mechanism.

Overall Results

For the 16 users who listen to more than 10 songs, each created a mean of 2.3 stations and listen to a mean of 40.5 minutes of music on each station.

Most users indicated that they liked using MegsRadio (mean 4.3 on scale from 1-5) and would likely recommend it to their friends (mean 3.9). Of the individual who indicated that they "might" or "probably" recommend the system, many of the specific comments focused on minor technical issues such as system did not work on a certain web browsers or that the music corpus (of 15K songs) was too limited. Some of the overall comments about MegsRadio were "It was like Pandora with improvements … it helps keep the genre and tastes of the listener in front" and "It definitely focuses the radio a lot. More music I want, less I don't."

Artist Similarity and Tag Controls

From a musicological perspective, we were interested in learning whether users preferred to use artist similarity, semantic tags, or both to control the radio. To assess this, we tracked the number of times an artist or tag was added to a station. We also asked users how "useful" they found these features in the post-study survey

Feature	Implicit Usage	Perceived Usefulness
Artists	1.6	4.6
Anti-Artists	0.5	4.1
Tags	0.4	4.3
Anti-Tags	0.1	4.2

Table 2. Comparison of using Artists and Tags to Control a Station. *Implicit Usage* describes the number of Artists and Tags that were added to each station on average. The mean *Perceived Usefulness* as determined by a post-study survey (on a scale from 1-5).

Table 2 shows the results of both our implicit usage statistics and the perceived usefulness as determined by users on the post-study survey. Artist-tags were used the most and were perceived as the most useful. This is not surprising since this is the most common way to seed stations on popular radio players. Tags and Anti-Artists (i.e., negatively associated artists) we also commonly used to control the radio station. Anti-tags used less frequently but this may be due to the fact that the anti-tag feature was not explicitly explained in our user interface. The perceived usefulness suggests that we should make it easier to use anti-tags.

About 24% of the stations were built using at least one artist and at least one tag. This is interesting to note because none of the personalized Internet radio players we have

surveyed allow users to create a station using both tag and artist seeds.

Station Characteristics

We were also interested in comparing the usefulness of our other control mechanisms. Table 3 compares the four song characteristic sliders: tempo, danceability, loudness and energy. While all four were well used, the energy slider was the most popular in both usage and perceived usefulness.

Feature	Implicit Usage	Perceived Usefulness
Tempo	0.4	3.9
Danceability	0.7	3.9
Loudness	0.6	3.8
Energy	1.1	4.5

Table 3. Comparison of Song Characteristic Sliders. Implicit Usage refers to how often each slider was manipulated during our study. Perceived usefulness refers to how users rate the usefulness of the control mechanism on a scale from 1-5.

Finally, we compare the four playlist options in Table 4. Users were most interested in the "Local Music" feature of MegsRadio though many turned this feature down when using the system. This may indicate that the playlist algorithm has been too aggressive at playing music by local musicians especially when the semantic relevance to the seed artists and tags was low relative to music by non-local artists.

Feature	Implicit Usage (More / Less)	Perceived Usefulness
Local Music	16% / 46%	3.9
Popular Music	27% / 19%	3.6
Obscure Music	13% / 13%	3.4
Repeated Songs	5% / 45%	3.4

Table 4. Comparison Playlist Characteristics. Implicit Usage refers to the percentage of stations that were left either indicating more, neutral or less of the type of music at the end of the study. Perceived usefulness refers to how users rate the usefulness of the control mechanism on a scale from 1-5.

Feedback

We asked a number of questions about the visual and auditory feedback. The most notable finding is that mean perceived usefulness of showing the upcoming artists was 4.6 out of 5. We also found that 89% of the users noticed that their choices directly affect the music that was being played. All of the users indicated that they trusted that their choices were being used to find relevant songs for them.

One user wrote, "I liked that there was visible feedback to know that when I was tweaking settings on what I wanted to listen to that it was actually happening." Another stated that she "loved watching [the upcoming artists] change without interrupting playback."

DISCUSSION

In this paper, we introduced MegsRadio as a tool for exploring a number of different control and feedback mechanisms. We found that both artist similarity and tags (genres, emotions, instruments, etc.) are useful for controlling the stream of music. Furthermore, we found that a significant number of users use artists and tags together despite the fact that we are unaware of any current radio player that offers this functionality. We also found that many other control features were useful for music discovery such as filtering by song energy and highlighting music by local artists.

We found that both auditory and visual feedback was perceived as important. Additionally, we found that showing upcoming artists provides the user with context for their interactions with the system. Many users indicated that seeing the upcoming artists change when they made a change to the station gave them better understanding of the playlist algorithm.

Finally, we note that the results from our small-scale user study are just a starting point for future exploration. Future usability studies will focus on A/B comparison that better isolate specific interactive design features (e.g., including vs. excluding tags for seeding stations). Finally, we are in the process of developing a new interface that incorporates much of the feedback that we received from our user test. This will hopefully allow us to better emphasize some of the features (e.g., local music recommendation) that our test subjects were enthusiastic about and indicated as being most useful.

REFERENCES

1. P. Lamere, *Tutorial on Using Visualization for Music Discovery*, International Society for Music Information Retrieval, 2009