
Using Personalized Radio to Enhance Local Music Discovery

Douglas R. Turnbull
Ithaca College
Ithaca NY, 14850
dturnbull@ithaca.edu

Justin A. Zupnick
Cornell University
Ithaca NY 14853
jzup90@gmail.com

Kristofer B. Stensland
Ithaca College
Ithaca NY, 14850
kstensl1@ithaca.edu

Andrew R. Horwitz
McGill University
Montreal, Quebec, CA
andrew.horwitz@mail.mcgill.ca

Alexander J. Wolf
Ithaca College
Ithaca NY, 14850
awolf1@ithaca.edu

Alexander E. Spirgel
Ithaca College
Ithaca NY, 14850
aspirge1@ithaca.edu

Stephen P. Meyerhofer
Ithaca College
Ithaca NY, 14850
smeyerh1@ithaca.edu

Thorsten Joachims
Cornell University
Ithaca NY, 14853
tj@cs.cornell.edu

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
Copyright is held by the author/owner(s).
CHI 2014, April 26–May 1, 2014, Toronto, Ontario, Canada.
ACM 978-1-4503-2474-8/14/04.
<http://dx.doi.org/10.1145/2559206.2581246>

Abstract

We explore the use of personalized radio to facilitate the discovery of music created by local artists. We describe a system called *MegsRadio.fm* that produces a customizable stream of music by both local and well-known (non-local) artists based on seed artists, tags, venues and/or location. We hypothesize that the more popular artists provide context for introducing new music by more obscure local artists. We also suggest that both the easy-to-use and serendipitous nature of the radio model are advantageous when designing a system to help individuals discover new music. Finally, we describe an interactive map that features personalized event recommendations based on the user's listening history. Results from a small-scale user study indicate that users are more aware of the local music scene after using it, discover relevant local music events, and would recommend the experience to others.

Author Keywords

Local Music Discovery; Personalized Internet Radio; Recommender Systems;

ACM Classification Keywords

H.5. Information interfaces and presentation (e.g., HCI): Miscellaneous. H.5.1 Multimedia Information Systems, H.5.5 Sound and Music Computing

7 Types of Digital Music Services

Artist Profile Sites: host music content for amateur and professional artists (e.g., Bandcamp, SoundCloud)

Celestial Jukeboxes: provide music consumers with on-demand access to millions of songs (e.g., Apple iTunes, Amazon MP3, Spotify)

Internet Radio: broadcast streams of music that are generated by human DJs or automatic playlist algorithms (e.g., Shoutcast, iHeartRadio)

Personalized Internet Radio: users can seed a station with an artist or tag and then control a customizable stream of music through a rating and feedback system (e.g., Pandora, Slacker, Jango)

Continued on next page...

Introduction

This work explores the use of personalized Internet radio as a means to develop local communities of artists and music listeners. Unlike commercial systems like Pandora, Last.fm, and Apple iTunes that tend to make conservative recommendations by mainstream artists [2], our aim is to create a professional-grade music recommendation system that is not subject to popular market forces, but that focuses on the growth of the local community and the societal, artistic and economic benefits that come with that growth.

To achieve this goal, our system called *MegsRadio.fm* provides a platform where local artists can showcase their work, and music listeners can discover local artists and events. In this paper, we describe the side of the system that faces the music listener. Similar to Pandora and Slacker, a music listener can seed station with a favorite artist or semantic tag (e.g., a genre, an emotion) and then listen to a customizable stream of related songs. However, our system plays a mixture of songs by more well-known artists alongside songs by local artists. The idea is that, at first, the songs by popular artists produce an enjoyable listening experience, as well as provide trust in the quality of the music recommendations. However, over time, songs by local artists will become familiar and enjoyable as well. That is, as music listeners, we tend to enjoy music that is familiar to us [7] so we gradually introduce new music in a contextualized manner. Finally, our system recommends local music events to each listener based on his or her listening habits and preferences.

The user study presented in this manuscript provides an initial datapoint that further informs the design of our and similar systems that aim to create awareness of local artistic communities. We consider *local* as being

any place that an artist would self-identify as having a connection to. This might include a place where the group formed, currently resides, or performs on a regular basis.

Topology of Digital Music Services

There are a number of technologies that can be used to aid the music discovery process. For the sake of clarity, we roughly divide them into seven groups based on their primary functional purpose (see sidebar.)

Much like the brick-and-mortar music retailers, the large celestial jukeboxes and (personalized) Internet radio players tend to promote music by popular artists. They also provide access to millions of songs by more obscure artists though the large majority of these songs are rarely purchased or listened to [3].

Local artists often use *artist profile sites* like Bandcamp and SoundCloud to host music content. While these sites provide limited search functionality, they are not generally designed to facilitate the discovery of novel artists within a local music community. *Event aggregators* like BandsInTown provide a listing of music events for a given area. They focus on regional and national touring bands and generally have sparse coverage of events outside of large metropolitan areas.

There have been a few efforts at improving event discovery at the hyperlocal level [4]. Frestyl [2] focuses on local music event recommendation in Berlin but does not offer users the ability to listen to music from local artists. Other sites like ListenLocalFirst (Washington DC), Austin Music Map (Austin TX), and Spindio (New York City) provide detailed music information and song samples for individual cities. Another new music service, DeliRadio, shows event listings and allows

7 Types of Digital Music Services (Cont.)

Music Encyclopedias:

contextual information (biographies, discographies, social tags, album cover artwork) about artists and their music (e.g., AllMusic, last.fm)

Music Blogs: detailed music reviews by experts and passionate fans (e.g., The Hype Machine)

Event Aggregators: provide musicians with a platform to alert local music fans when they are playing in the area (e.g., BandsInTown, Songkick)

users to listen to artists playing at nearby venues. Despite its name, DeliRadio operates more like a celestial jukebox in that users select specific songs in an on-demand manner.

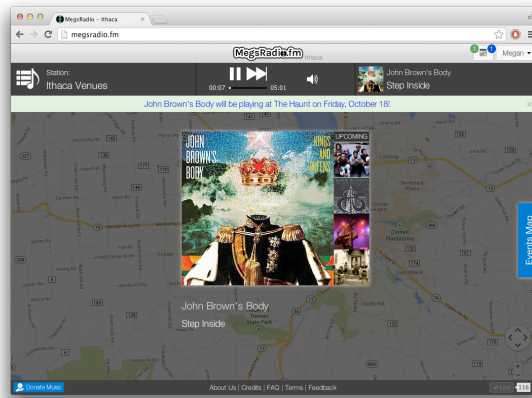


Figure 1: Main Player - Users can create and modify customizable radio stations. They are also notified if the current song is by a local artist or if the artist has an upcoming show at a nearby venue.

MegsRadio.fm: Designing for Local Music Discovery

Like DeliRadio, we designed MegsRadio.fm to help users discover music from a local geographic region. As opposed to their celestial jukebox model, we explore the personalized radio model for a number of reasons. First, radio listeners do not know which songs will be played next, providing an element of serendipity. By contrast with the celestial jukebox model, users create playlists from a collection of known songs.

Second, personalized radio players require relatively less work to use when compared with on-demand

systems. Whereas on-demand systems generally require users to select each song for a playlist, personalized radio players only require one seed artist or tag before a back end playlist algorithm creates an endless stream of music.

Third, the radio model allows us to better contextualize music by local artists. For example, when a user seeds a station with a popular artist, the recommendation algorithm creates a selection of songs by both popular (non-local) artists and local artists. The more well-known songs ground the station with relevant and familiar favorites, while discovery happens by slowly introducing songs by similar-sounding local artists.

Compared with other personalized radio players like Pandora, MegsRadio.fm incorporates a number of unique features that focus the user's attention on local music discovery. For example, in addition to creating a station with an artist or semantic tag (e.g., genre, emotion, instrument), MegsRadio.fm users can start stations for special tags that are related to cities, venues, or festivals. We also make it easy for users to create arbitrarily complex stations by seeding a station with multiple artists and tags. For example, a user can listen to a station with songs that sound similar to "The Beatles", are "aggressive", and might be heard at the "South by Southwest" music festival. In addition, users can control the mix of music along a variety of characteristics such as the ratio of local to non-local songs or popular to obscure songs (see figure 2). Users can also filter songs on a variety of acoustic properties such as tempo, energy, and danceability [5].

Another unique feature is an interactive map that ties the user's listening experience to a list of events in the region. Users can filter events by artists that they have

liked or listened to while using the radio player. Events can be sorted by date or relevance where relevance is based on a score derived from the user's listening history and song preference ratings.

User Study

Ithaca, a small city of about 30,000 permanent residents in central New York, has been selected as the testing ground for MegsRadio.fm since it is where the authors currently reside. Ithaca has both a high concentration of musicians, events and festivals, and is home to two large academic institutions (Cornell University and Ithaca College) with a total of approximately 27,000 students.

Participants were recruited from a local music festival, Craigslist, and through the researchers' social networks. No close friends or family members were recruited for this study. 34 participants (24 male, 10 female, avg. age = 25.88, $SD = 7.04$) completed a pre-survey about their music consumption habits and involvement with the local music scene. We then asked them to use MegsRadio.fm whenever they would normally be listening to music on their personal computers.

Of the 34 who completed the pre-survey, 19 completed a post-survey and only 11 were local residents of Ithaca. The post-survey asked about their level of MegsRadio.fm usage, perceptions of the application in comparison to other existing music services, and measured their awareness and engagement with the local music scene in Ithaca.

Feedback

Of our 11 test subjects who were from Ithaca and completed both surveys, 9 indicated that they would continue to use MegsRadio.fm for music discovery and 10 indicated that they would recommend the system to others. Some users stated that they preferred using MegsRadio.fm to existing commercial systems such as Pandora. For example, one stated, "MegsRadio was

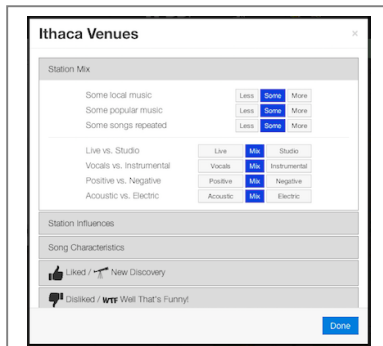


Figure 2 - Station Preference Menu - gives the user control over the mix between local vs. non-local artists, popular vs. obscure artists, and a number of other acoustic characteristics. Users can add and remove *station influences* (not shown), which can include other artists, genres, emotions, instruments, venues, festivals, and locations.

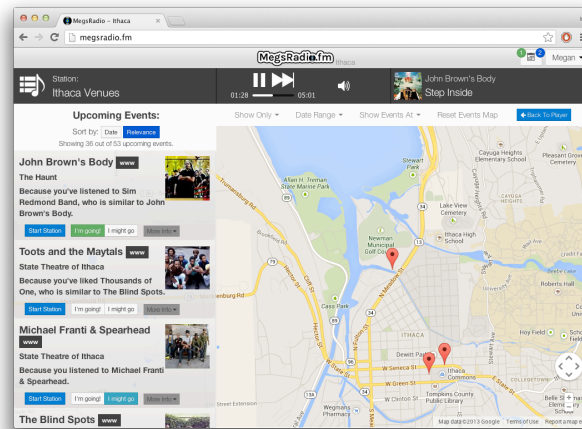


Figure 3: Events Map - displays upcoming events as both a list and as markers on a map. Users can search for events by location, date, listening history or liked artists. Events can be sorted by date or relevance.

Both backend playlist generation and event recommendation algorithms make use of a large music index [9] that relates songs and artists to tags, events, and other similar artists. The index is populated with information that is collected from various web APIs (e.g., Echo Nest, Last.fm), harvested from websites (e.g., local events calendars), and automatically derived from the audio content using machine listening [6, 9]. Our use of machine listening is particularly important for this application since many of the local artists tend to have little (if any) contextual information available through online sources.

easy to navigate and gave me a lot more control over building my playlist based on narrow parameters than other services. I particularly liked the ability to guide the playlist by increasing/decreasing the influence of certain artists, tags, etc. This made it easy for my stations to learn to match my tastes."

7 of the 11 test subjects indicated an increase in "awareness of the local music scene" and only 1 indicated a decrease based on difference in their responses between the pre-survey and post-survey. When asked to comment on distinctive features of MegsRadio.fm, nearly everyone pointed to the connection with local music recommendation. For example, one said that MegsRadio.fm was *"Much more interesting in terms of localizing the music, confining your listening to people who do/may play in local venues in the near future."*

One interesting finding is that 8 test subjects indicated a decrease on their level of agreement with the statement that they "attend music events on a regular basis". However, of these 8 individuals, 6 stated that they were likely to attend one or more events that they had discovered while using MegsRadio.fm. One possible explanation for this is that as the user's perception of the size of the local music scene increases, the perception of his or her own involvement within that scene decreases. This hypothesis is something we plan to explore in future studies.

The most positive feedback seemed to focus on the events map (see figure 3) and event recommendation. For example, a typical comment was that the *"map feature was unique and different than I've seen on any other sites."* One user comment in particular captured our overall design focus: *"hearing is better than a*

"description" of the band. Descriptions are subjective and may not be accurate; if you hear it, you KNOW if you approve or not. Two bands in particular I never thought about, and once I heard them, I would now be very inclined to go to their show."

Negative feedback tended to focus on technical challenges that we will need to address. Some users experienced slow page load times and buggy interface behavior. Others noticed issues with our data, such as one artist having multiple names (e.g., Björk vs. Bjork). This resulted in playlists that oversampled music from these duplicated artists. Users also complained that our music corpus (currently about 25K tracks) seemed small relative to other commercial systems (> 1M tracks).

Implications of Current Work

Our initial intention was for users to discover music using the personalized radio player in a passive manner and then become interested in attending related music events. However, a number of users seemed to focus first on the events map and then launch stations to learn about local artists in a more active manner.

This active approach to music discovery may be related to our finding that while users reported an increase in awareness of the local music scene, they also reported a decrease in their perceived level of attendance at music events. This may be due to users' increased perceptions of the size of the music scene, rather than a decrease in actual attendance. There lies the possibility that the larger virtual representation of the music scene imparted its way into the users' own mental representations of the music scene. This was an unexpected finding, which might lend more credence to the concept of *fear of missing out*, a previously studied

attitude held by social media users [8]. Exploiting this fear of missing out (e.g., encouraging users not to miss out on relevant shows) within a constructive community-building platform like MegsRadio.fm may prove to be a positive design practice.

Future Work

Based on the feedback from our pilot study, we have some early indications that our approach to using personalized radio can increase awareness of and participation in the local music scene. In the coming months, we plan to launch MegsRadio.fm as a free service (ad-free, non-profit) for local music discovery in other communities beyond Ithaca. Once we attract a stable base of users, we will be better able to explore some of the hypotheses we have presented in the paper. Specifically, we plan to measure whether or not users attend more events over time after using MegsRadio.fm on a regular basis.

We are currently developing a mobile app that will be tied into our existing backend recommendation system and act as a functional equivalent to our web app. We are also developing a companion site that artists can use to upload music, add upcoming events, create a biography (with location-specific information), and attach (or correct existing) tags to their music. This system will provide us with a secondary source of human feedback (in addition to preference ratings from listeners using MegsRadio.fm) and will be useful for improving our machine listening and web scraping algorithms. Artists can also use this companion site to both receive analytical information about and communicate with their listeners. Finally, we have begun working with local venues and music festivals to create special location- and event-specific tags so that

This work is supported by NSF Grants IIS-1217485 and IIS-1217686.

users can become familiar with artists and songs before events happen.

References

1. Bassoli, A., & Brewer, J. frestyl: simplifying the process of promoting and discovering local live music. In *Proc. ACM SIGCHI Italian Chapter 2011*, ACM Press (2013), pp. 167-170.
2. Celma, O. *Music Recommendation and discovery: Long Tail, Long Fail, and Long Play in the Digital Music Space*. Springer, (2010). pp. 139-145.
3. Elberse, A. *Blockbusters: Hit-Making, Risk-Taking, And The Big Business Of Entertainment*. Henry Holt and Co. (2013)
4. Hu, Y., Farnham, S. D., & Monroy-Hernández, A. Whoooly: Facilitating information seeking for hyperlocal communities using social media. In *Proc. CHI 2013*, ACM Press (2013), pp. 3481-3490.
5. Jehan T. & DesRoches D., The Echo Nest Analyzer Documentation. http://developer.echonest.com/docs/v4/_static/AnalyzeDocumentation.pdf (2011)
6. McFee, B., Barrington, L., & Lanckriet, G., Learning content similarity for music recommendation. *Audio, Speech, and Language Processing, IEEE Transactions on*, 20(8), (2012) pp. 2207-2218.
7. North, A. C., & Hargreaves, D. J. Subjective complexity, familiarity, and liking for popular music. *Psychomusicology: Music, Mind & Brain*, 14(1), (2008) pp. 77-93.
8. Przybylski, A. K., Murayama, K., DeHaan, C. R., & Gladwell, V. (2013). Motivational, emotional, and behavioral correlates of fear of missing out. *Computers in Human Behavior*, 29(4), 1841-1848.
9. Turnbull, D. Indexing Music with Tags. *Music Data Mining*, CRC Press, (2011), pp. 276-302.