Chapter 7

Promoting Social Interaction in Public Spaces: The Flytrap Active Environment

Andrew Crossen and Jay Budzik

Northwestern University Intelligent Information Laboratory

1 Introduction

Flytrap is an active environment that knows its users' musical tastes and can automatically construct a soundtrack that tries to please everyone in the room [5]. The system works by paying attention to the music that people listen to on their computers. Users of the system have radio frequency ID badges that let the system know when they are nearby. Using the preference information it has gathered from watching its users, and knowledge of how genres of music interrelate, how artists have influenced each other, and what kinds of transitions between songs people tend to make, the "virtual DJ" finds a compromise and chooses a song. The system tries to satisfy the tastes of people in the room, but it also makes a play list that fits its own notion of what should come next. Once it has chosen a song, music is automatically broadcast and played.

Music lives at the boundary of private and public life, and a discussion about music preference can provide a great deal of insight about the people involved in the discussion. Explanations for music preference are often personal—inaccessible to those who do not share a certain social context, set of interests, or intimate knowledge of the listener. As such, social interaction involving music preference provides a context in which the boundary between public and private can be explored. Listening to music in a public setting can provide the basis for social interaction about typically private matters. Expectations about the listener are drawn from the music to which they are listening. Yet because of the boundaries implicit in music as a representation of self—that not everyone knows the back-story of a certain artist or genre, nor have they trained themselves to decipher lyrics obscured by style—music allows the listener to maintain a boundary around aspects of their private life that can be inferred from their musical taste, in a way that invites a predefined class of others in, and actively keeps others out.

Technologies that can make personal preferences visible in a public setting provide opportunities for studying these boundaries and the effects of manipulating them. We focus here on a technology that intends to spur social interaction in a public setting by making personal preferences visible to those present in the space. Of interest to us are technologies that can facilitate informal social interaction by highlighting common ground. The technology we discuss represents a new kind of social environment that actively facilitates participation in social exchange by manipulating representations of self. Musical tracks that match people's preferences replace the generic, "lowest common denominator" music typically played in public environments. In exchange, users delegate control over sharing their preferences to our technology. Having technology mediate this sharing allows us to explore the issues that arise around privacy and trust.

We built the Flytrap system so we could explore these ideas.

2 Motivation

People's choice of music can be deeply personal, as seen in Tia DeNora's ethnographic studies of music listeners. Listening to music often triggers memories of events and experiences that have emotional significance. Explanations for music preference often involve very intimate, private matters. One interview subject noted that she often privately listened to Schubert's *Impromptus* because they "reminded her of her father" and listening relaxed her before work [6].

Yet the very personal experience of listening to music originated in a completely public setting. Before modern listening technologies, music was performed at small social gatherings in homes and public places, as well as in larger venues that are more like today's concerts. Even with the widespread use of portable, private listening devices today, the act of listening to music often occurs openly in public. It is this dichotomy between the private and personal nature of the *reasons for listening* to certain music—and the public settings in which the *act of listening* to music occurs—that makes music preference particularly compelling.

Listening to music in public is often an invitation to the public to enter the private space of the listener. Readers will be reminded of times they heard music come from someone's office, stopped in to ask what it was, and heard a very personal story about why the listener liked the music they were playing so much. Readers will also be reminded of times when an explanation for the music choice was avoided, perhaps because the reason was just too personal.

This ability for the listener to invite people in, but still negotiate the boundary between public and private by choosing what, if anything, to share about their personal reasons for choosing certain music, is one of the properties of music preference that make it particularly well suited for this study. People ultimately control whether the reasons for liking certain music is shared, even if they choose to publicly expose what music they like.

Music often also has properties that allow the listener to select a portion of the public to invite in, and a portion to keep out. These built-in boundaries are present from opera (which typically filters in the wealthy and educated and filters out the rest) to punk rock or hip-hop (which do the opposite). Readers might remember hearing music coming from an office, and deciding not to stop in, because there was no interest in what was being played. At the same time, commonalities in music preference (especially if that preference isn't shared by many people, generally) often pre-qualify people for certain kinds of social exchanges.

Erving Goffman suggests that focused interaction in an encounter becomes possible when commonality is perceived among the participants [8]. The most interesting commonalities are those experiences and preferences that aren't shared by many others. This *rare common ground* provides the Flytrap active environment the opportunity to promote social interaction among its users.

2.1 Public space as a locus for social interaction

Many public spaces exist mainly to support specific kinds of interactions. Convention halls, classrooms, and train stations are all designed to support specific uses. The constraints of the activities performed in these spaces govern the social interactions taking place therein, which are often focused around a task. Music does not typically play a part in the rules of social interaction in these spaces, and could be construed as a distraction in some cases (e.g., if it was playing during a lecture).

Other spaces are designed for public use but do not typically facilitate or result in social interaction. Elevators, waiting rooms and subway cars are sometimes even designed to make social interaction entirely optional, if not difficult (it is difficult to talk with someone who is reading a magazine in a waiting room). People are barely socially present in these spaces.

Still other spaces—company lunch rooms, pubs, parks, and green space—are designed to provide a venue for the public to interact socially on a more informal basis. Music in these spaces establishes the mood and social constraints of an informal environment [6], and helps people to mold a socially appropriate dramaturgical “front” to fit the space [9]. Some of these spaces position music as a focal point, while others use it as a backdrop to other activities. Coffee houses may play soft music as part of establishing an intimate atmosphere, while rock clubs play loud music that make quiet interaction difficult.

These types of spaces—where people are socially present in an informal manner—best suit a Flytrap installation. These spaces balance informality, social accessibility, and generality of use for our augmented environments. Flytrap's synthesis of personal preferences for use in public can help to support and ultimately change the character of these public spaces, using music as a backdrop to or basis for social interaction.[1]

[1] The properties and character of public space outlined above can be used to tailor the kind of music played. This idea is further explored in Section 6.

2.2 Can we do better than elevator music?

While it is feasible to suggest that most people visiting a blues club like blues, in many settings it is impossible to infer music preference from the location alone. Acquiring and synthesizing knowledge of the tastes of people present in these types of general-purpose public spaces is a difficult task. Often the perceived difficulty outweighs the perceived benefit and “elevator” music is played, or music is omitted altogether. Ideally, the musical tastes of each individual present and knowledge of the activities they intend to perform could be used to come up with better mix of music to play.

According to Joseph Lanza [11], during the early 1980s, patrons of Pittsburgh's airport complained of feeling uneasy waiting for their flights as Brian Eno’s ambient composition “Music for Airports” played over the public address system. This was a particularly poor choice of background music given that people made specific complaints against it, instead of ignoring it altogether. We see an opportunity to fill generic public spaces with something more interesting than “elevator” music. We want to present visitors in the space not just with something they like, but with a selection of their music that could promote sharing and social interaction by weaving a musical thread amongst those in the room.

Our goal with Flytrap is to leverage rare common ground in music preference to provide a basis for social interaction. Flytrap invites people to share the intimate reasons for their choice of music with each other as they feel comfortable. Public spaces can thereby be transformed into environments that facilitate new kinds of social interaction.

3 Related Work

Already, environments are designed to support and encourage certain types of behaviors. Yet the environments of today are often static, unless someone is actively orchestrating the experience of those within them. The environments of the future—the rooms, offices, churches, shopping malls, and parks of the world—will know about the people inside of them, and will be able to craft an experience for those people that not only reflects their preferences, but actively supports everything they do, including encouraging social interaction and discourse, when appropriate.

One active environment is in use today at a company health club. The MusicFX system [13] is a group music preference arbitration system installed in a company fitness center. Users of the gym sign up for MusicFX by detailing their music preferences in a survey. Users rate about 100 genre-constrained radio channels on a five-point scale, used later by a voting mechanism to choose the radio channel to play based on the preferences of those present in the gym at a given time. As their music preferences change, users of MusicFX update their survey profile.

At first glance, the goals and functionality of MusicFX and Flytrap seem similar. Both systems strive to replace “lowest common denominator” music in a public space by democratizing the music selection process. Both systems model user music preferences to make recommendations in a group context. Yet the systems differ in their intent: while MusicFX aims to make the environment less offensive and more enjoyable, Flytrap aims to make it socially engaging. The difference in choice of deployment environment reflects this difference in intent.

Flytrap can foster social interaction by exposing specific overlapping preferences that are not generally shared. In order to do so, Flytrap's information about user preferences must be highly granular (the system must know about specific artists and songs). Because of the volume of preference information required, we focused on implicit sources of music preference that can be gathered automatically, without requiring explicit user input.

This difference in focus and approach resulted in significantly different design choices, reflected below.

4 The Flytrap System

A number of distributed components comprise the functionality of the Flytrap system as a whole. Each user has a Flytrap agent tied to their personal media player, responsible for gathering information about their music preferences, and voting on songs being considered for play in a group setting. A central server houses a database of song information populated by each personal agent, and a file repository containing each of the musical tracks. Public areas where the system is to operate are outfitted with a voting agent, user identification subsystem and music player.

4.1.1 Gathering music preferences

Deriving a user's musical tastes from observation provides a more accurate characterization of the user's tastes than a survey, which requires somewhat difficult introspection and exhaustive enumeration. Likewise, observation allows the system to gather preferences in context, which provides fertile ground for research on more context-sensitive methods.

Because of the personal and spontaneous nature of music selection, it is necessary to capture people's listening habits in an unobtrusive way, participating directly in the act as a silent observer. Many users generally listen to music on one of a handful of media players on their personal computers. We developed interfaces to several popular Windows-based media. These give us access to the music a user is currently interested in. While preferences are not gathered from portable personal music devices, the model extends to such devices if they support the installation and use of third-party software.

On each Flytrap user's personal machine, their instrumented media player gathers information about what tracks the user is listening to, and records this in Flytrap's database. Users interact with their media player just as they always do, with no additional work on their part to make music known to the Flytrap system. The tracks themselves are uploaded to the server by the agent if they don't exist in the central repository. In this manner, the user's preferences are learned through the act of listening to the music as they would normally.

4.1.2 Approach to music representation and recommendation

There are generally two approaches to building recommender systems: statistical approaches, which mine usage or preference data in order to provide recommendations (e.g., collaborative filtering [16], or market-basket analysis [7]); and content-based approaches, which use knowledge of items and similarity among them to provide recommendations (e.g., FindMe systems [1, 2]). The primary advantage of statistical approaches is they require no representation of items and no theory of similarity among them. Their disadvantage is that they require large volumes of rating data before they begin to make recommendations that make sense to users. Likewise, they require an item to be rated in order for it to be recommended. Content- or knowledge-based approaches are the opposite: they require knowledge of items and a theory of similarity, yet they require no rating information in order to make recommendations.

We chose to implement the following content-based recommendation algorithm given:

1. We wanted our system to introduce users to music they hadn't heard before
2. The number of artists available would be large and initially unrated
3. A large knowledge-base about music and genres was immediately available
4. We needed to compute similarity in order to maintain play list coherence (see below)

The system did not immediately meet the requirements of typical statistical methods, and therefore lent itself naturally to the application of content-based recommendation frameworks.

As a basis for recommendation, Flytrap needs to know information about each musical track it has been exposed to via the personal agents. Flytrap uses multiple methods of obtaining information about a piece of music. One technique involves looking at the metadata in MP3 files (called ID3 tags) to determine a track's artist and genre information. Since the genres reported in ID3 tags are notoriously inconsistent, we use a web wrapper ([4]) built around the AllMusic Guide[2], a popular music information site, as a backup to retrieve the genre of the track given the artist. This covers cases in which the genre is available but the artist isn't. In cases in which no information is available about the track, it is not included in the track database. This is a minor failing that, in practice, does not occur very often. However, statistical recommendation algorithms could be leveraged to cover instances when tracks are not properly tagged, or artists are simply not known to the system. Ideally, recommendations would not be based on loose concepts like genre that shift over time. Artists can change genres, and tracks can vary wildly in genre

[2] Available at http://www.allmusic.com/

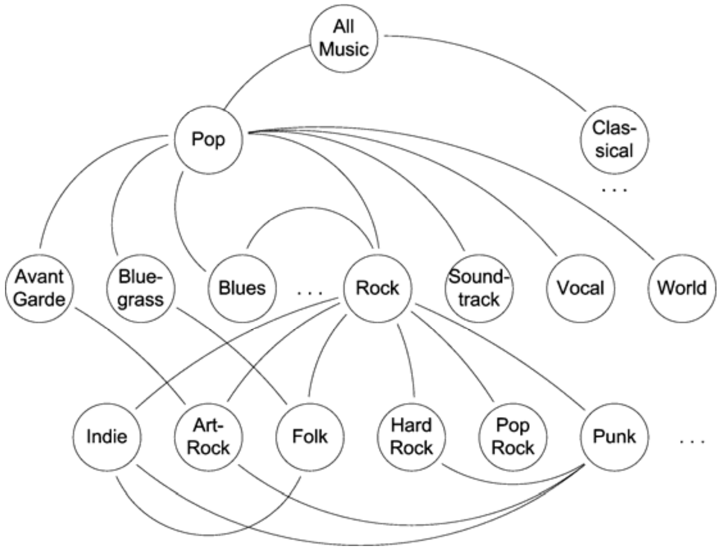

Figure 1: Partial Music Genre Network. This highly connected network of genres is derived from a crawl of Allmusic.com, an online music guide. This semantic network forms the basis of our content-based recommendation algorithm.

and tone, even on the same album. A representation that captured track-level information that reflects this reality would improve the quality of the play lists our system constructs. This could be accomplished by analyzing the acoustical properties of the tracks themselves, and is the subject of future work. Notwithstanding the above shortcomings, our system produces results that are generally good enough for its users.

A song's genre acts as a key into a semantic network of inter-related genres, which are used to determine similarity among artists. A similarity network of genres derived from the AllMusic Guide is used to determine similarity between genres (see Figure 1). Artists can belong to multiple genres. Links between genres are assigned weights based on the number of artists they have in common. This similarity information is used in a variety of ways to compute a group play list, as described below.

4.1.3 Determining who is in a public space

As users move from their personal spaces into the Flytrap-equipped public space, the system needs to know who is present in the location. We used a system of radio frequency ID badges (see [18,19]) that transmit a unique identifier to a base station. Each badge's unique identification number is tied to a user's Flytrap music profile. As users walk into the public area, their unique ID is picked up by the base station, informing the system that a user is now present in the space.

4.1.4 Deciding what tracks to play

After personal preferences are gathered and the system notices a user has entered a Flytrap-equipped public space, the system begins the recommendation process. Flytrap decides the next track to add to a play list using a voting mechanism whereby the agents representing each user present in the room give a numerical vote to each track in the system's database each time a new track event is signaled. The criteria for voting are based on artist, genre and style information, as follows:

- A user's Flytrap agent will give a song a high vote if it's an artist they've listened to previously, and a higher vote to those they've listened to frequently.
- Songs users present have never listened to before receive positive votes from the user's agent if the genre is the same or similar by some degree to music they'd previously listened to.

Similarity among artists (and thereby songs) is computed by spreading activation (see [15]) along the links in the genre network. For each user, the activation level for a given genre is the ratio of the number of tracks they've listened to over a certain time in that genre to the number of tracks they've ever listened to in that same time frame. Thus, if N is the number of times a user has played any track during a given time period (counting a track twice, for example, if it was played twice in that period), and $N(G)$ is the number of times the user has played a track in genre G. Then the activation level of genre G, $A(G)$, is given by:

$$A(G) = N(G) / N \tag{1}$$

The activation level of an adjacent genre G' is the ratio of the activation level of G and the number of links out of G. So if $|\mathrm{adj}(G)|$ is the number of nodes connected to G, $A(G')$ is given by:

$$A(G') = A(G) / |\mathrm{adj}(G)| \tag{2}$$

The value of a user's vote for a given artist is given by the sum of the activation levels of the genres to which that artist belongs multiplied by the percentage of users that would have voted for that song. The more frequently a song is preferred across all users in the database, the less likely it will be played. This ensures the system exposes the rare common ground that exists among the users participating in the experience it crafts.

Once the voting has completed, the sum of each agent's votes (normalized by the number of users) form a probability distribution across the entire database of songs. Songs that get more votes have a higher probability of being played. Songs that get few votes can still be played, but it's less likely. We chose a stochastic algorithm so the system could be somewhat serendipitous, causing users to become aware of new kinds of music.

In addition to the personal user agents, the system also has a disc jockey (DJ) agent, which has the power to override and manipulate the outcome the voting process based on its own 'good taste'. The rules followed by the DJ agent are much like those a human DJ would use in deciding what to play next:

- Never play two tracks by the same artist in a row.
- Maintain loose genre coherence across tracks.

Unless it's "Two-fer Tuesday" on a radio station, a human DJ will not typically play the same artist twice in a row. The Flytrap DJ agent assigns very low probabilities to songs by artists whose songs were played the last 10 times. The result is a less repetitious play list, and also one that frequently drifts into new areas, because this rule significantly reduces the number of choices in a given genre available for play.

In order to produce play sequences with as few jolting transitions as possible (e.g., playing hard rock after classical), the DJ agent uses its similarity network of genres to assign new probabilities to each track, based on the candidate track's genre and the genre of the track it just played. The probability associated with each candidate track is multiplied by its genre similarity to the previous track, as captured by the semantic network of genres described above. As a result, the DJ will favor new tracks from the same (or similar) genres as the track that was just played. The result is a new probability distribution over the entire database of tracks, which the DJ uses to choose the next song.

4.1.5 Playing the music

Once all votes are cast and a song is selected, the winning track is streamed across the network for play on the machine located in the public space. Music that the system votes on is housed in a central, network-addressable repository to facilitate streaming to the playback machine.

In sum, the system understands the music its users like and broadcasts that music in the spaces it controls.

5 Experiences and Iterations

We installed Flytrap in one of our public areas used for demonstrations, informal student lunches, and studying. This space perfectly suited a Flytrap installation. Its physical design and utility promote informal activity, and its inhabitants are often people that do not know each other intimately. Graduate students and professors regularly use the area for various purposes.

We gave ID badges to around a dozen people and outfitted their personal music players with the Flytrap preference gathering agent. Over 3000 musical tracks were listened to by users during that period (and therefore made available to Flytrap). For several months, we kept the system running so people could give us feedback. Through a series of informal interviews and observations, we derived the following improvements and iterations, which we implemented and deployed.

5.1 Promoting New Music

One behavior observed as a result of the initial voting mechanism was that users were not being exposed to enough new music. Preferred artists would consistently receive higher votes and the system would oscillate between the same few artists. We added a rule to the DJ agent that selects music users in the space hadn't heard yet, using the same spreading activation recommendation model responsible for selecting songs (except that it punishes songs users heard recently). This provides necessary noise in the system to ensure the play list doesn't get "stuck" in a poorly-connected sub-graph of the genre and artist network.

5.2 The Vote Visualizer

After about a week of users interacting with the system, they began to ask us why particular songs were chosen. Flytrap maintained an internal model of its song selection rationale, but offered no visual representation of the process. From this feedback we built a vote visualization component (see Figure 2).

The vote visualizer graphically depicts the voting process in real time. Each user is assigned a color when their badge is first picked up by the system. Candidate track titles have a text color based on an interpolation between the user's color and the strength of their vote on the track. Brighter track graphics represent stronger votes. As votes are tallied, the track names meander around the screen. Those with higher weights gravitate toward the center, and then the DJ's vote is calculated and the final track selection is made and highlighted (the song's final position lies on a circle with radius proportional to the probability of that song being played). This gives the user not only a sense of how the voting process is going and a visual cue as to the winning track, but also shows the outliers–those tracks that lost but were also strong candidates, as well as some of those that weren't.

Figure 2: The first version of the Flytrap vote visualizer, working with two users. The first user is a Bob Dylan fan, and Flytrap plays the Dylan cut "Tell Me, Momma" (top-left). The second user, a Johnny Cash fan, enters, and Flytrap plays a track from a Cash & Dylan album, "T for Texas" (bottom-right). Songs near the center are more likely to be played.

Users reacted positively. For those in our Flytrap-enabled space, this visual reinforcement of how their private musical preferences overlap promoted interesting conversation.

While the system was having its intended effect, users sometimes voiced concerns over whether putting their preferences up on the screen in front of everyone was an invasion of their privacy. Although we intended to get some push-back around privacy issues, we didn't expect users from the same group to feel uncomfortable openly sharing their preferences with each other. The visualization also wasn't scalable: with more than two or three people in the room, colors got muddied, and there was too much text on the screen to read it.

A second iteration on the vote visualizer can be seen in Figure 3. This version makes the association between users and music less visible, balancing privacy concerns with the goal of providing enough common ground to spur interaction. It also promotes additional insight into the voting process (which was the original user request). This iteration shows users the intermediate recommendation sets built by the system during the system's process of deciding what to play next, and by explaining its choice in limited English. The revised interface also scales much better to dozens of users and organizes information in a way that is easier to consume.

Figure 3: Redesigned interface for the Flytrap system. This version presents the system's current and future behavior without disclosing the preferences of individual users. Instead, users can share these preferences with each other only if they so choose.

The "Top Picks of People Present" lists artists ranked highest by the agents representing those present, without disclosing their identities. Instead, the ranking and number of people who influenced that choice are presented. Since the system has no representation for the content of a given track, track names were not displayed, because they didn't convey any useful information. This change was intended to directly address the privacy concerns voiced by users.

The "Next Up" list is the virtual DJ's working set of recommendations. As described above, this list contains any direct overlap in users' artist preferences, as can be seen in the selection of Wilco and Shins. Music preferred by one user directly but not another may be added if the degree of genre similarity is high enough (Cash). New music not directly preferred by users (i.e., not in the list of top picks) is also added by the DJ based on the degree of genre overlap between the candidate artist and artists in user profiles. In the example above, The Pogues and Billy Bragg are added because of the high degree of preference for the music's genre expressed by users present in the space (in the example above, Wilco, Crosby, Stills, Nash and Young, and Johnny Cash influence the choice of The Pogues and Billy Bragg).

"Now Playing" represents the current selection being played. Full artist, song and album information are displayed. In addition, the system explains why it chose the song it's playing, which includes a combination of user preferences (genre preference or artist preference) and similarity to the last song played.

At a glance a user can see what the system is doing and why. Displaying a selection of the preferences of users present allow those people to start conversations about any of the artists mentioned. Not displaying which users influenced what choices allows users to choose whether they share that information.

6 Future Work

During the time Flytrap was running, users suggested new features they would have liked included in the system. In addition, we noticed opportunities for improvement ourselves based on an analysis of the system's behavior. The following future work reflects this.

6.1 Reflections on Flytrap's Preference Model

The preference model employed by Flytrap represents choices of individuals for themselves and uses those choices in contexts in which others are present. The implicit assumption is that the relationships among the participants in the experience crafted by Flytrap are irrelevant. This assumption is reasonable, given the music ultimately selected by the system is influenced by all of the participants equally. However, it could benefit the system to understand the context of its presentation more deeply (e.g., that there is a meeting in the room and people are waiting for it to begin, vs. it's lunchtime). Choices made in private are not always the same as those made in public [10].

6.1.1 The use of space and social context of inhabitants

DeNora suggests that "music is active in defining situations because, like all devices or technologies, it is often linked, through convention, to social scenarios, often according to the social uses for which it was initially produced ..." [6]. Environments are often used for activities other than their central purpose. The classroom that by day supports student learning in a traditional structured manner is also cleaned at night, by a completely different set of people.

Flytrap should be sensitive to the use of a space. One way to establish such context is to use properties of the music itself to determine the appropriateness of a recommendation. Muzak[3] engineers listening experiences based on the stimulus level of different pieces of music. The stimulus level is based on, among other factors, tempo and instrumentation. AllMusic[4] maintains a wealth of descriptive

[3] Available at http://www.muzak.com/
[4] Available at http://www.allmusic.com/

data about artists, including the style and mood of their music.

Additional tags including the tempo, instrumentation, and lyrical content could be used to functionally describe social contexts in which the music should be played. Music with no lyrics, a slow to moderate tempo and low volume might be appropriate play during a meeting in a conference room. That same room during lunchtime might be better served with an up-tempo number at a higher volume. In addition, Flytrap could integrate with calendar systems to become more aware of how certain spaces are planned to be used, and when.

Moreover, the social contexts in which groups of people are engaged can differ dramatically moment to moment. An informal hallway chat can turn instantly formal when the boss walks up. Users should be able to more directly influence the sets of their music played in certain places. Flytrap's representation of its users should be expanded by better understanding the relationships among them. This could be done, for example, by using the employee LDAP directory or social networks like Friendster[5].

6.2 Richer Music Representation

Flytrap's representation of music ends at a genre-level. It has no representation of a track (or individual song) other than what artist performed it, and therefore what genre it belongs to. More granular information about the music it is playing could dramatically improve the character of the play lists it constructs.

6.2.1 Music Content Analysis

Research into analysis of the content of musical waveforms will lead to a richer model of recommendation based on the sound of the song. DJs regularly "beat match" music in clubs to provide smooth transitions between songs of different speeds. Songs can be slowed down or sped up without altering the pitch of the music through algorithms (see [17]) designed to pinpoint and adjust the rhythmic content of music. Untapped aspects of recommendation for Flytrap involve looking at appropriate times for playing fast or slow music, and assessing a personal music collection for trends in rhythmic structure to provide more on-point recommendations in the future.

Moreover, technologies that allow the programmatic detection of melody in a rich, complex waveform can be leveraged to derive an even richer content model.

6.2.2 Environmental Annotations

Some modern media players give users a view into their music collection that is based on the time of day in which they tend to play various tacks. This is a stepping

[5] Available at http://www.friendster.com/

stone into richer environmental contexts associated with a music collection. We mocked up a component of Flytrap that records the time of day, weather, and season for each piece of music played. Over time, user profiles can be built to describe patterns in their listening habits. These patterns can be taken into account when Flytrap makes recommendations in a group setting.

A simple example is a user who tends to listen to Christmas music around the holidays and no other time of year. While this music is in their collection, it should not be included in the set of possible recommendations in July. Another example is a user who tends to listen to Billie Holiday on dark rainy mornings. In a group setting where the weather is dark and rainy and vocal jazz is a centroid of recommendation in the current social context, it should be weighted higher. These kinds of patterns are indicative of some users' listening habits and could be used by Flytrap to make better recommendations.

6.3 Personal Annotations

Music is often internally catalogued by how it came into one's life. Nick Hornby's book *High Fidelity* involves a protagonist who organizes his music collection by the ex-girlfriend that introduced it to him. This organizational mnemonic reminded that character about the time when the music was fresh to him, and invoked subsequent recall of other life events at that time.

Music comes into peoples' lives for many reasons and from many sources. It may be from browsing the shelves at a music store, hearing an opening act at a concert, or hearing it emanate from a car window. Because we can only capture the behavior of a user in context of them using their personal media player—thus capturing the moment when music enters a personal store—it is feasible to outfit the Flytrap user agent with functionality that lets users add personal annotations to their music.

Radio stations often have a call-in request line that lets listeners select and dedicate songs. The radio station DJ empowers the listener by giving them an opportunity to establish and publicly broadcast an intimate personal connection to the music. The Flytrap user agent could be instrumented to let a user associate media with each artist or song. These media could then be uploaded to the central repository and associated with the artist or song and user in Flytrap's database. When the song is next selected for play by Flytrap in a public space, the associated media could be presented.

These additional representations add a more personal dimension to music played by Flytrap in a group setting. To other listeners, these personal associations signify that the music about to be played is especially representative of that individual in the manner the associated media portrays. The window into that user's private preferences is opened a little further, giving others additional context for making personal connections to that person. Aside from just being fun, this is an interesting means of attaching and recalling very personal musical artifacts in a way that can be used by Flytrap to promote social awareness of these connections.

6.4 User Vote Control

The Jukola system [14] offers patrons control of music at a local bar with a remote control that lets users override automatic selection. Flytrap is currently an autonomous system that leaves the final song choice up to the Virtual DJ for subsequent broadcasting. Users of our system expressed an interest in being able to influence the choices once they'd been made. We imagine a Flytrap ID badge outfitted with "thumbs up" and "thumbs down" buttons. As a new song is selected and played by Flytrap, users can express their like or dislike of the choice by pressing one of the respective buttons. The user's preference is transmitted to Flytrap and used in future voting decisions concerning that user.

The feedback mechanism is ambiguous in giving the user only two controls with which to express an opinion. By providing negative feedback, did the user mean they disliked the artist or was it just the particular song? Could it have been the genre of music? A set of rules that drive learning user feedback preferences over time fine tunes this process:

– If a user provides negative feedback for a song, a new recommendation is generated immediately and an association is made between that track and the user who expressed their dislike.
– If that user gives negative feedback again for a different song by the same artist, that artist will no longer be played when that user is present in the space.
– If a song by the same artist is played and the user is present but provides no feedback, the assumption is that they disliked the particular song. That song will never be played when the user is present.
– If a user provides positive feedback for a song, that song is weighted stronger in the future.
– If a user provides positive feedback for another song by the same artist, that artist is weighted stronger in general for that user.

This type of interaction lets users feel more in the loop about their influence on the music being played. Behind the scenes, a profile of the idiosyncrasies of a user's music preferences is built for future research into more context sensitive methods of recommendation. However, our intent is to have the system's selections delight its users initially, lessening the importance of this feedback mechanism.

7 Conclusion

Flytrap was built with the purpose of exploring how personal preferences could be gathered and manipulated to craft a group experience in a public setting. Flytrap is a work in progress, yet it provides an example of a new kind of system: one that deliberately manipulates boundaries to achieve social effects. Though the system has not been widely tested, our informal evaluation showed such environments can be built, and that they can achieve those ends. Along with more extensive deployment and evaluation, we envision our next steps leading us closer to a general model of leveraging personal preference in a group setting.

We see Flytrap as a mediator of music and associated social interactions in all environments and group contexts. Imagine the use of Flytrap spreading to waiting rooms and building lobbies. Opportunities for exploiting common ground through music are possible in each of these spaces. Willis promotes the notion that music's powers are best seen in action [20]. Flytrap manages music intelligently in a space populated by people with disparate tastes. Discovering music that exploits commonalities amongst the occupants of a space establishes a basis for social interaction. Interactions occur at the ground level of social action, exposing personal preferences in a manner conducive to sharing.

DeNora suggests in *Music in Everyday Life* that "music can be used … as a resource for making sense of situations, as something of which people may become aware when they are trying to determine or tune an ongoing situation." By providing music that brings together the personal preferences of everyone in a space, rare common ground can be highlighted to produce an active, social environment.

Acknowledgements

The authors would like to thank Kris Hammond and Larry Birnbaum for supporting this work, and Don Norman for helpful discussions about it. Andrei Scheinkman and Sanjay Sood contributed to the ideas presented here. We would like also to thank Louis Lapat for helpful discussions about this topic connected with his work on *Pregunte a Latino Chicago*, a separate music project.

References

1. Burke, R., Hammond, K. & Cooper, E. Knowledge-based navigation of complex information spaces. In *Proceedings of the 13th National Conference on Artificial Intelligence*, pages 462-468, AAAI, 1996.

2. Burke, Robin. The Wasabi Personal Shopper: a case-based recommender system. In *Proceedings of the sixteenth national conference on Artificial Intelligence and the eleventh Innovative Applications of Artificial Intelligence conference*, pages 844-849, AAAI, 1999.

3. Churchill, E., Girgensohn, A., Nelson, L., and Lee, A. "Blending Digital and Physical Spaces for Ubiquitous Community Participation." In *Communications of the ACM*, February 2004.

4. Crossen, A., Budzik, J., and Hammond, K. J. Between Now and the Semantic Web. In *Proceedings of the ICAPS 2004 Workshop on Automated Planning and Scheduling*. 2004.

5. Crossen, A., Budzik, J., and Hammond, K. J. Flytrap: Intelligent Group Music Recommendation. In *Proceedings of Intelligent User Interfaces 2002*. ACM Press, 2002.

6. DeNora, Tia. *Music in Everyday Life*. Cambridge: Cambridge University Press, 2000.

7. Fu, X., Budzik, J., and Hammond, K. J. Mining Navigation History for Recommendation. In *Proceedings of Intelligent User Interfaces 2000*. ACM Press, 2000.

8. Goffman, Erving. *Encounters. Two studies in the sociology of interaction*. Bobbs-Merrill: Indianapolis, 1961.

9. Goffman, Erving. *The Presentation of Self in Everyday Life*. Doubleday: Garden City, New York, 1959.

10. Hebdige, Dick. *Subculture: The Meaning of Style*. Methuen: New York, 1979.

11. Lanza, Joseph. *Elevator Music*. University of Michigan Press, 2004.

12. Feiner, S., MacIntyre, B., and Seligmann, D., "Knowledge-based augmented reality". Communications of the ACM 1993. 36(7), pp. 52-62, July 1993.

13. McCarthy, J. F., and. Anagnost. T. D. (1998). MUSICFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts. In *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work* (CSCW '98), Seattle, pp. 363-372. 1998.

14. O'Hara, K., Lipson, M., Jansen, M., Unger, A., Jeffries, H., and Macer, P. Jukola: democratic music choice in a public space. In *Proceedings of the 2004 conference on Designing Interactive Systems* (DIS 2004), ACM Press, 2004.

15. Quillian, M. Semantic Memory in *Semantic Information Processing*, M. Minsky (ed.), MIT Press, 1968.

16. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J., GroupLens: An Open Architecture for Collaborative Filtering of Netnews, In *Proceedings of ACM Conference on Computer Supported Cooperative Work*, 1994.

17. Sethares, W. A., Morris, R. D., and Sethares, J. C., Beat tracking of audio signals using low level audio features, IEEE Trans. On Speech and Audio Processing. Vol. 13, No. 2, March 2005.

18. TIRIS. Texas Instruments Radio Frequency Identification (RFID). See http://www.ti.com/tiris/ (Accessed October, 2004).

19. Want, R., Hopper, A., Falcao, V., and Gibbons, J. "The Active Badge Location System". In *ACM Transactions on Information Systems*, 10(1):91-102, January 1992.

20. Willis, Paul E. *Profane Culture*. Routledge & K. Paul, 1978.