# Flytrap: Intelligent Group Music Recommendation

**Andrew Crossen, Jay Budzik, and Kristian J. Hammond**
Intelligent Information Laboratory
Northwestern University
1890 Maple Ave., Evanston, IL 60201 USA
+1 847 467-1265
{crossen, budzik, hammond}@infolab.northwestern.edu

## ABSTRACT
Flytrap is a group music environment that knows its users' musical tastes and can automatically construct a soundtrack that tries to please everyone in the room. The system works by paying attention to what music 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 playlist that fits its own notion of what should come next. Once it has chosen a song, music is automatically broadcast over the network and played on the closest machine.

## Keywords
Intelligent environments, ubiquitous computing, audio spaces.

## 1. INTRODUCTION
We describe work on managing music in an environment (e.g, [1], [3], [5]) shared by multiple people. Traditionally music in shared environments is selected based on what will offend the least people, and please as many as possible. Acquiring and synthesizing knowledge of the tastes of people present is difficult. 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 a more accurate notion of what to play. Flytrap combines presence awareness with profiling in an attempt to address this. The system uses an Active badge system ([4]) to uniquely identify a badge wearer. Flytrap watches users play music on their personal computer and records the tracks as well as information about the music to a database. Flytrap uses this track data and general knowledge of how genres of music interrelate to make decisions about what a group of users might want to listen to.

### 1.1 Example Interaction
Flytrap is installed on a machine in our public demonstration area. When no one is around, the system is dormant. A lab member with his RF ID badge walks into the demo area, where the proximity badge reader picks up his presence. This user is a Sammy Hagar fan, so Flytrap plays a track from Hagar's first album. Another user, who listens mainly to Mozart, enters the space. Flytrap fades out the current track, and picks a new one. The system comes up with a song that neither of the users know, yet they both enjoy; a track from Brian Eno the system has captured from another lab member. Flytrap decided to play Eno because his music lies between Mozart and Hagar. Ambient was heavily influenced by classical. Eno's music is classified as both ambient and progressive rock. Progressive rock is a sub-genre of arena rock, which is how Hagar's music is classified.


*IUI'02*, January 13-16, 2002, San Francisco, California, USA.


## 2. FLYTRAP ARCHITECTURE
Each user has a Flytrap agent, responsible for gathering information about their music preferences, and voting on songs being considered for play in a group setting.

### 2.1 Track Information Recorder
On each Flytrap user's personal machine, their RealJukebox (available at http://www.real.com/) music player is instrumented with software that gathers information about what tracks the user is listening to, and records this in Flytrap's database. The tracks themselves are uploaded to the server 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, unlike other systems (e.g [5]), which require the user to fill out a survey. 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).

### 2.2 Genre Network
Flytrap uses metadata in MP3 files (called ID3 tags) to determine artist and genre information about tracks. A song's genre acts as a key into a semantic network of inter-related genres, which are used to determine similarity among artists. Since the genres reported in ID3 tags are notoriously bad, we use a web wrapper ([2]) built around AllMusic, a popular music information site (available at http://www.allmusic.com/), to retrieve the genre of the track given the artist. A hand-built similarity network of genres is used to determine similarity between genres. This similarity information is used in a variety of ways to compute a group playlist, described below. Each link in the network is labeled by a similarity score between 0 and 1. We scored relations between 200 of the most popular genres for this initial system.

### 2.3 Voting Mechanism
Flytrap decides the next track to add to a playlist using a voting mechanism whereby an agent representing each user present in the room gives a numerical vote to each track in the system's database each time a new track event is signaled. The criteria for voting are

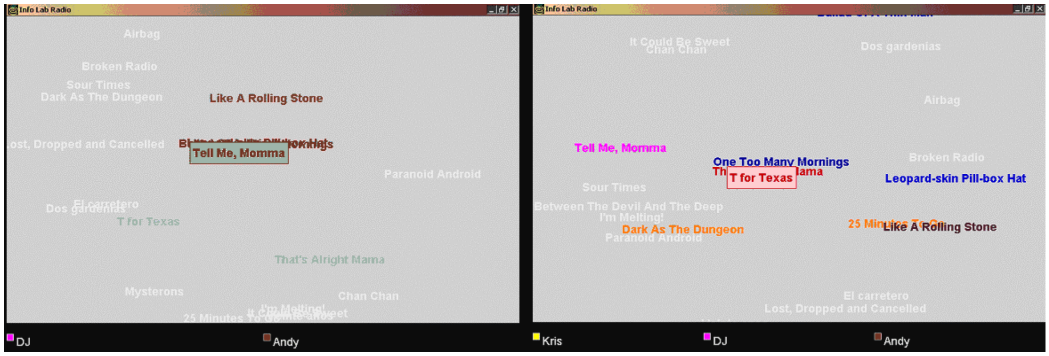


**Fig. 1. Visualizer results for two users. The first is a Bob Dylan fan, and Flytrap plays the Dylan cut 'Tell Me, Momma' (left). The second user, a Johnny Cash fan, enters, and Flytrap plays a track from a Cash & Dylan album, 'T for Texas' (right). Songs near the center are more likely to be played.**

based on artist and genre information. A user's Flytrap agent will give a song a high vote if it's an artist they've listened to previously. Songs the user has 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. Once the voting has completed, the agents' votes are combined and then normalized into 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 representation so the system could be somewhat serendipitous, causing users to become aware of new kinds of music.

In addition to human user agents, the system also has a disc jockey agent, which has the power to override and prioritize the outcome of the user's vote 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: (1) never play two tracks by the same artist in a row and (2) maintain loose genre coherence across tracks. Unless it's 'Two-fer Tuesday' on a radio station, a DJ will not typically play the same artist twice in a row. The 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.

## 2.4 Vote Visualizer

To show users of the system how the votes are turning out, 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 (their final position lies on a circle with radius proportional to their probability). 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 those that weren't.

# 3. FUTURE WORK

The first order of enhancements to Flytrap will be to incorporate more context cues into track selection. Having the track information recorder gather personal statistics about the frequency of play of particular artists and genres, and what was listened to before and after a given track could further hone the selection process. Wave analysis and comparison of tracks could also provide further insight. Environmental cues such as the time of day the music is listened to, active applications on the user's computer, and what the weather is outside can provide deeper and more interesting mood-task-music correlations. Also of interest is adding a notion of 'space ownership' so that the longer a user stays in the space, the more influence they have. We see this as being particularly useful for spaces like offices. Broader use and user evaluation will also help us make additions and changes to Flytrap. To this end, we hope to someday set up a dance club on campus that uses this technology.

# 4. CONCLUSION

Flytrap manages music intelligently in a space populated by people with disparate tastes. It allows you to discover new music in the system by playing tracks by artists who you haven't listened to before but have qualities similar to music that you enjoy. The system also provides the potential to meet new people in this manner; two people down the hall from each other may listen to similar music and not be aware of it, but their similar interests may be brought to bear when using Flytrap.

# 5. REFERENCES

[1] Coen, M., 'Design Principles for Intelligent Environments,' in *Proceedings of AAAI '98*1998.

[2] Crossen, A., Budzik, J., and Hammond, K., 'XLibris: An Automated Library Research Assistant,' in *Proc.of IUI 2001*, 2001.

[3] Franklin, D., 'Cooperating with people: The Intelligent Classroom,' in *Proc. of AAAI '98,* 1998.

[4] Hopper, A., Harter, A., and Blackie, T., 'The Active Badge System,' in *Proc. of INTERCHI '93*, 1993.

[5] McCarthy, J., and Anagnost, T., 'MusicFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts,' in *Proc. of CSCW '98*, 1998.