

## Online Game Feeds Music Search Engine Project at UC San Diego

*Toward a music search engine that lets you type in regular words and returns songs*

September 26, 2007

Daniel Kane

Gert Lanckriet, an Electrical and Computer Engineering Professor at UCSD, is leading an effort to create a natural-language search engine for music.

Listen Game is an online, multiplayer game prototype that gets players to label songs with words.

UC San Diego electrical engineers and computer scientists are working together on a computerized system that will make it easy for people who are not music experts (like the senior author's mom) to find the kind of music they want to listen to - without knowing the names of artists or songs.

In a new paper, the researchers demonstrate that the online music game they created provides data crucial for building the back-end of a music search engine that allows users to type in words in order to find songs.

"When my mom gets up in the morning and is like, 'I need some energy to go jogging,' she has no clue what title or artist is going to help her with that," said Gert Lanckriet, the UCSD electrical engineering professor overseeing the project.

What Gert's mom needs is a "Google for music" - a search engine for music that lets you type in regular words like "high energy instrumental with piano," "funky guitar solos" or "upbeat music with female vocals," and get songs in return.

One option for creating this kind of natural-language music search engine is to manually annotate as many songs as possible - but this is expensive and limits the depth and breadth of the search engine's reach. Another option is to train computers to do the song annotations.

The UCSD researchers have, in fact, built such a system over the last two years. They call it a computer audition system. You feed it songs and it annotates them, thanks to a series of algorithms they created. Once a song is annotated, you can retrieve it using a text-based search engine. But before the system can start annotating songs, it has to be trained - via a process of machine learning. Getting enough data to properly train the system to label a wide range of music accurately is difficult.

At an academic conference on music information retrieval this week in Vienna, Austria, the UCSD researchers describe important progress toward this goal. Lanckriet and others from UCSD report that an online music game they created, called Listen Game, is capable of capturing the crucial word-song combinations that are needed to train their system to label large numbers of songs automatically.

Listen Game is an online, multiplayer game prototype that gets players to label songs with words. Like the popular image labeling games, such as ESP game, Peekaboom and Phetch, Listen Game relies on people surfing the Web to generate valuable data while playing the game. For Listen Game, the "human computation"

occurs when users listen to clips of songs and determine which words are most and least relevant to the song. Users earn points when they pick the same words as others who are playing at the same time and listening to the same song clips.

Collections of reliable word-song combinations are of great interest to the UCSD researchers. They use these "semantic annotations of music" to train their home-grown computer system to label songs it has not previously encountered.

In their paper, the UCSD researchers show that the word-song combinations generated by people playing Listen Game can be used to accurately train their system to label new songs.

"We've shown - in academic terms - that our game works. We're close to the performance we get with comparable survey data from the music undergrads we paid to fill out music surveys," said Doug Turnbull, a computer science Ph.D. student at UCSD who, along with Lanckriet, is an author on three papers presented at ISMIR, the music information retrieval conference.

Doug Turnbull, a computer science Ph.D. student at UCSD is combining his love for music with computer science.

Thanks to a grant from the UCSD Jacobs School of Engineering's von Liebig Center, the team is working with interactive designers to make the game more fun and engaging. With the new games, the researchers hope to collect the data necessary to continue to develop their automated music annotation system.

"A bunch of engineers made Listen Game. The second generation will have a livelier look and feel," said Lanckriet.

The researchers will release new games later in 2007. The games will be free and will carry no advertising. Players will be able to play anonymously or they can log in and customize and personalize the games and socialize with other players in real time. Send an email to [gert@ece.ucsd.edu](mailto:gert@ece.ucsd.edu) to be notified when the new games are released.

In addition to the ISMIR paper focused on Listen Game, the UCSD researchers are presenting a paper at the same conference on identifying words that are most likely to be meaningful to music search engines and another paper on detecting boundaries within songs.

Work on the word identification paper began two years ago, when the researchers were mining text from music reviews of specific songs, in search of words that are meaningful and useful in the context of a music search engine.

"If you look at a music review, there are so many words that are not relevant, you want to filter them out to get the quality training data, to get words that are acoustically describing the song," said David Torres, a UCSD computer science student working on his master's degree and an author on the paper.

The authors propose an approach for identifying musically meaningful words and show that this kind of filtering can improve the performance of the music annotation and retrieval system.

They also highlight complications that arise from the fact that music is subjective. "For example, a pre-teen girl might consider a Backstreet Boys song to be 'touching and powerful' whereas a dj at an indie radio station may consider it 'abrasive and pathetic,'" the authors write. To account for this subjectivity, electrical engineering Ph.D. student and paper author Luke Barrington explained that the team considers the level of human agreement for different words when filtering through the music annotations they collect.

In the third paper, Turnbull, Lanckriet and authors from Japan's National Institute of Advanced Industrial Science and Technology, present their work on detecting boundaries between musical segments in a song, such as between a verse and a chorus. The researchers' strategy for automatic boundary detection could be useful for generating music thumbnails, for efficient music browsing and for music information retrieval.

"In order to analyze the audio of a song, it is useful to be able to chop it up into meaningful segments," said first author Doug Turnbull.

Media Contact: Daniel Kane, 858-534-3262

