

Engineers use song-annotating algorithms to study music playlists (w/ Video)

October 27 2009

**SMARTER THAN GENIUS?
HUMAN EVALUATION OF MUSIC RECOMMENDER SYSTEMS.**

Luke Barrington Reid Oda* Gert Lanckriet
 Electrical & Computer Engineering, Cognitive Science
 University of California, San Diego



Figure 4. Genius ignores context. When subjects were unfamiliar with a seed song that they liked and had no information about song and artist names (20 trials), Genius matches the performance of the content-based system.

Electrical engineers recently pitted Genius -- the music recommendation system in Apple's iTunes -- against two experimental music recommender systems. Genius appears to capture acoustic similarities among songs within the same playlist, the researchers found. The University of California, San Diego electrical engineers also discovered that the music recommender they built from scratch can generate song playlists that human subjects thought were as good as those that Genius generates. The UC San Diego system works for songs that Genius knows nothing about. Credit: UC San Diego Jacobs School of Engineering

Electrical engineers recently pitted Genius - the music recommendation system in Apple's iTunes - against two experimental music recommender systems. Genius appears to capture acoustic similarities among songs within the same playlist, the researchers found. The University of California, San Diego electrical engineers also discovered that the music recommender they built from scratch can generate song playlists that human subjects thought were as good as those that Genius generates. The UC San Diego system works for songs that Genius knows nothing

about.

UC San Diego electrical engineering Ph.D. student Luke Barrington presented these findings on October 28 at the 2009 International Society for Music Information Retrieval Conference (ISMIR 2009) in Kobe, Japan.

"Our goal is to make a music recommendation tool that is as good as or better than Genius, but that does not require massive amounts of user data. The system we are developing can analyze and recommend completely unknown songs by new bands as accurately as it analyzes the most popular hits," said Barrington, who used the same underlying technology to create a series of music discovery games for Facebook and a new kind of music search engine that will be available for beta testing next week (<http://herdit.org/music/index.html>).

Tools for creating automated music playlists are increasingly useful now that huge numbers of songs are available for download and streaming to anyone with an Internet connection. iTunes - the most popular [music retailer](#) on the planet - has sold more than 6 billion tracks. Genius uses "collaborative filtering" on these purchase statistics to help people organize their music and discover new songs they might like based on similarity to a "seed" song that they do like.

By averaging statistics about how millions of listeners purchase and play music, Genius appears to actually capture acoustic similarities between songs, according to the new research, which involved human evaluation of music recommendation systems and was led by researchers at the UC San Diego Jacobs School of Engineering. Because Genius is a proprietary system whose secrets are not available to the public, the researchers studied it by testing its song recommendations against comparable song suggestions from experimental music recommender systems that they fully understood.

"At first we thought that collaborative filtering would just spot similar artists: if you like Daft Punk, it might recommend any song by Justice. However, we found that acoustic similarities come through in Genius playlists. Genius seems to find correspondences between songs that go beyond what you get by just matching artists. We don't think Genius actually knows anything about the acoustics of songs, but it still produces great recommendations since the data it is using is built upon acoustic analyses of the music done by millions humans," said Luke Barrington, an electrical engineering Ph.D. candidate at the UC San Diego Jacobs School of Engineering, and the first author of the new study. Gert Lanckriet, an electrical and computer engineering professor from UC San Diego and Reid Oda, an undergraduate in cognitive sciences at UC San Diego are co-authors on this paper.

Smarter than Genius?

The UC San Diego researchers found that the playlist generator they built using their own algorithms performed as well as Genius under certain conditions. In contrast to Genius, which uses information about the songs people buy and listen to in iTunes in order to learn which songs in any iTunes library are related, the UCSD music recommender relies on auto-tagging algorithms that use machine learning to label songs with descriptive words based only on the acoustic content of the songs.

"Our computer system works by listening to the music - it doesn't know anything about artists or albums or charts. In some trials of our survey, we tried to remove these biases from the human listeners by hiding the names of the songs and artists and making sure that subjects liked the seed songs but had not heard them before," said Barrington. In these cases, the researchers found that people liked playlists generated by the UCSD auto-tagging algorithms as often as they liked Genius playlists.

This partial parity with Genius underscores how the UCSD auto-tagging

algorithms can be used to generate high quality music playlists that incorporate lesser-known and unknown songs. Genius currently ignores relatively unknown songs because it lacks adequate wisdom from iTunes customers about how these songs connect to other songs. Systems like the auto-tagging music algorithms developed at UC San Diego could be useful in filling in the "blind spots" in Genius and other collaborative filtering systems that rely on the wisdom of the masses to generate playlists.

"We weren't expecting our system to beat Genius at making playlists based on the most popular songs - our system doesn't know about artists, popularity, release dates, albums or anything else that the average music fan is aware of. Once we add that information in, we think we can build something that is really smarter than Genius," said Barrington.

Music Games are Serious Research

The UCSD electrical engineers are always improving their auto-tagging algorithms thanks to a series of music discovery games they created for Facebook called Herd It. Players listen to songs and compete with other online players to most accurately describe the songs they are listening to. The electrical engineers use the song-word combinations these games generate to improve the accuracy and breadth of their auto-tagging algorithms.

As with the Herd It games, the surveys the engineers created to study playlists also began with humans listening to a song clip. Next, the survey system produced two competing five-song playlists for the same seed song. Playlists were generated by one of the four possible recommender systems the researchers were studying: Genius; Artist Similarity - a music recommender that selected songs based only on similar artists; Similar Tags - a music recommender that selected songs based only on the audio content of the songs (this system was built upon the UCSD

auto-tagging algorithms); or random playlist.

For each pair of playlists, subjects identified the list they thought was better. Playlists were built from the authors' personal music library of over 12,000 relatively popular songs that span the most common genres of Western popular music. The list of songs used to seed playlists is available at: <http://cosmal.ucsd.edu/cal/projects/playlist/>

People are moving more and more to streaming music, where listeners don't have the actual file. With access to every song on YouTube, deciding which of the 10 million songs you want to listen to next can be daunting.

"Playlists are an increasingly important tool for overcoming what otherwise might be an overwhelming amount of streaming music," said Barrington.

New Music Search Technologies

The algorithms that automatically tag songs with descriptive words that are at the foundation of the UC San Diego playlist generator are also powering a new breed of search engine for music as well as [Facebook](#) games that provide researchers with the information needed to improve the new search engine. To watch a video highlighting the capabilities of the new [music](#) search engine, go to: <http://cse-ece-ucsd.blogspot.com/2009/10/new-music-search-engine-on-way.html>

More information: Read the research paper at:

<http://cosmal.ucsd.edu/cal/pubs/Barrington-Genius-ISMIR09.pdf>

Source: University of California - San Diego ([news](#) : [web](#))

Citation: Engineers use song-annotating algorithms to study music playlists (w/ Video) (2009, October 27) retrieved 23 April 2024 from <https://phys.org/news/2009-10-song-annotating-algorithms-music-playlists-video.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.