

Applying Content-Based Recommendation to Indian Classical Music

Andrew Ash

October 20, 2008

Motivation

Every year, more music is produced than any single person can ever have the chance to hear. Sifting through all that music is an insurmountable challenge, so in the pursuit of a fresh sound many audiophiles spend more time listening to mediocrity than masterpieces. One of the largest obstacles to the growth of the music industry is the music discovery problem: what can be done to help consumers find the music they enjoy without having them subjected to everything else?

This problem falls in the domain of Music Information Retrieval, the science of computationally gathering useful information from music. One branch of the field is Content-Based Recommendation (CBR), which seeks to help solve the music discovery problem by producing suggestions based on a user's listening habits. More well-known approaches to recommendation include focuses on social patterns in musical taste (Last.fm) and expert-based content analysis (Pandora). The CBR approach aims to use the computationally-extracted musical contents, rather than a track's metadata, to group and recommend individual songs to its users.

Technique

We begin making a recommendation using the CBR technique by working off the assumption that if a listener enjoys a certain song, then another song with similar characteristics will be enjoyed as well. Then we find ways to quantify various attributes of the songs in the musical space, those which are dissimilar enough across the song space to have a discriminatory effect. These features are assembled into a feature vector representing that particular song, which is a point in n -dimensional space R^n .

In order to compare two vectors, we must choose a distance metric that determines the relative distance, or how "close" those points are, which is often chosen to be Euclidean distance. Using this distance relation, we can place the information for all pairings of songs in the space into a symmetric similarity matrix, with a cell's value as the correlation of the two songs represented in its row and column. Now we have all the information necessary to query the matrix with a search song and find the songs with the highest correlation coefficients.

These songs represent those which are the recommendations given that a user enjoys a certain seed song. Other, more detailed models can follow the history of a user's listening habits and modify the recommendation system based on that added information.

Objectives

We have assembled a team of five students at Georgia Tech in the Music Technology Group and Center for Music Technology to explore CBR in a specific domain: Indian Classical Music. Led by Professor Parag Chordia, the group aims to adapt current CBR models to take advantage of the intrinsic qualities of this musical domain, such as the emphasis on polyphonic melodies, rhythmic-focused pieces, and a more varied range of scales, or raags. As the recommendation system is developed, a subdomain of the project is to create the means of sharing these recommendations with the world. The roadmap for distributing the information is to create a "radio station" application on the Facebook app platform. This app will be the face of the CBR system to users around the world interested in discovering Indian Classical Music.

Team Role

My role in the team is presently focused on the storage and retrieval of the results of the recommendation algorithm. With n songs, there are $\binom{n}{2}$ distance metrics that will need to be stored and retrieved in an efficient manner. As the size of the music library increases, the storage requirements for the distance matrix increase on the order of $O(n^2)$. Besides the storage complexities of holding such quantities of data, significant work will also need to be done exploring the best way to take a seed song and user feedback to generate radio playlists. I will be exploring current literature regarding recommendation systems and collaborating with other team members and my mentor professor to create and fine-tune the recommendation algorithm.

Personal Background

My experience with Music Technology has been primarily through a class I took in the Spring of 2008 with Professor Chordia, MUSI 3450 Survey of Music Technology. One of the class assignments was a naïve implementation of a pattern matching system, using the spectral centroid of two types of tabla strokes. The techniques used in this approach will be extended and cross-applied to the work being done on this Indian CBR project. In addition to this experience, my Computer Science major will assist greatly in the implementation of the storage and retrieval portions of the project. As the preparations for beginning this research have continued, my previously general interests in algorithmically analyzing music have been focused on Content-Based Recommendations, and I look forward to this research.