

Applying Content-Based Recommendation to Indian Music

Andrew Ash

February 25, 2009

With the overwhelming availability of music online, listeners have few ways to discover new music tailored to their tastes. In this content-based recommendation project, we look at creating recommendations for listeners based on automated analysis of songs' musical attributes, similar in goal to Last.fm and Pandora (though they use social trending and manual annotation approaches instead). By analyzing the spectral features of many short snippets of sound, we can construct statistical representations for individual songs, appearing as ellipsoids in high dimensional spaces described by Gaussian mixture models (GMMs). Comparing two songs is then performed by calculating their representations' distance via Earth Mover's Distance, and recommendations are created by selecting nearby songs.

The novelty of this project surfaces in its customization of the feature set to Indian music. By taking advantage of this musical subset's characteristic attributes, particularly raag, monophony, and tempo, we hope to produce improved recommendations on our target body of roughly 5,000 tracks. To measure improvement, we both rely on standard evaluation measures such as genre and artist R-precision and look to listener feedback from our web interface and to our intuition of the recommendations' reasonableness. With a baseline now established on solely spectral features (MFCCs), we will move to measuring the effect of adding temporal and melodic features to the recommendation system. We expect to soon see performance improvement, particularly as we increase our track count and begin experimentally determining the most discriminative features of that musical body.

Research Team: Meghu Adoni, Andrew Ash, Oliver Jan, and Karan Mehra
Led by: Dr. Parag Chordia