

A PREFERENCE RANKING MODEL USING A DISCRIMINATIVELY-TRAINED CLASSIFIER

Jeremy Reed

School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332
jeremy.reed@gatech.edu

Chin-Hui Lee

School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332
chl@ece.gatech.edu

ABSTRACT

This paper presents an ordinal regression approach to the query-by-description problem. Instead of returning a single classification, such as genre, or a list of the top N songs assumed to be relevant, this algorithm mirrors choices similar to "like", "skip", "play", and "hate" buttons seen on commercial Internet radio stations. Ordinal regression can be viewed as a hybrid between multi-class classification and regression. While ordinal regression defines a small number of discrete classes similar to multi-class classification, there is a dimension of preference which determines the order of classes, which is similar to regression. To model a given song, a transcription of acoustic segment models is derived from an unsupervised approach using hidden Markov models. These transcriptions serve as text documents for latent semantic analysis. A discriminatively-trained ordinal regression algorithm provides the ranks for a given song. Test queries and ground-truth labels are developed using textual musical descriptions from songs to combat the problem that listeners do not base musical similarity solely on acoustic perception. Initial results on a small set of queries indicate the proposed ordinal regression approach performs better than a common perceptron-based ordinal regression algorithm. Further, the results using acoustic segment models is shown to be superior compared to a typical bag-of-frames algorithm using Gaussian mixture models. Extensions to real playlists and user preference indicators are currently in development.