

Abstract

Of late, the interdisciplinary science of music information retrieval (MIR) has seen a gained interest in automatic playlist generation. Selecting the “suitable” songs that blend well with our mood and putting them in the “correct” order can be achieved by hand, but not always are we right in stringing them appropriately into a playlist. Automatically generating playlist replaces this by learning the song information and seeks to produce playlists consonant with one’s mood.

This project builds on these endeavours by introducing Au2seed, an automatic playlist generator that uses a modified Gaussian process prior learned from albums, user playlists, similar songs and user listening history. Gaussian process regression is carried out to learn a user preference function over songs, which is influenced by a user’s mood. The final covariance kernel takes song metadata as inputs. We unveiled 10 different combinations of song features to be used as inputs, leading us to 10 different covariance kernels, which gave rise to 10 different *versions* of Au2seed. An assessment of these playlist versions has shown that they perform better than mere random generation of songs. Concurrently, we selected a handful of song features that are important in song similarity, and identified those that are dispensable. Au2seed v3.6, which we deemed the best out of the 10 versions, has displayed reasonable capabilities qualitatively when we put it to human evaluation and experimented with it for ourselves.

Keywords: music information retrieval, gaussian process regression, playlist generation, meta-training