

UNSUPERVISED LEARNING OF MELODIC SEGMENTATION: A MEMORY-BASED APPROACH

Miguel Ferrand, Peter Nelson
University of Edinburgh, UK

Geraint Wiggins
City University, London, UK

Background

Music segmentation is recognised as an important step in the abstraction of musical contents and researchers have attempted to explain and formalise how listeners identify such constituent units. Several attempts to model music segmentation have employed Gestalt principles to identify discontinuities in a musical surface, and it is often suggested that some of these principles operate independently of the listeners' musical knowledge. Other theories suggest that listeners memorise recurrent features in the music and use this knowledge to carry out musical analytical tasks. Leonard Meyer emphasised the relationship between learning and expectation, where the unpredictability of an event can change its prominence and hence constitute an auditory marker within a sequence of events.

Aims

We conjecture that some of the perceived features responsible for the identification of segment boundaries in melodic sequences can be learned from a corpus of non-annotated musical data. Given this hypothesis, the aim of this research is to develop a computational model that can simulate this learning process and then use the acquired information to predict the location of segment boundaries for a given melody.

Method

The implementation of the proposed model is based on a mixed-memory Markov model. The model is constructed from the regularities found in parametric sequences (pitch and duration based) corresponding to all the melodies in a given training set. Then, for a given target melody, sequence probabilities are estimated from the model and prediction entropy profiles obtained. Further analysis of these profiles reveals distinct intervals of entropy change, between which, we conjecture, segment boundaries are likely to be found.

Results

An on-going study on melodic segmentation performed with listeners, provides segmentation data for a set of melody examples. Preliminary results seem to indicate that the boundaries generated by the model are located in the vicinity of the boundaries indicated by the listeners.

Conclusions

We suggest that salient intra and inter-textual features can be learned from a musical corpus in an unsupervised fashion. We then demonstrate how the acquired information can be used to predict the location of segmentation boundaries in melodies, without any a priori knowledge.