Learning about structural analysis from structural analyses

Jordan B. L. Smith (1)*, Ching-Hua Chuan (2), Elaine Chew (1)

(1) Queen Mary University of London. London, UK
(2) University of North Florida. Jacksonville, FL, USA
* = Jordan.Smith@eecs.qmul.ac.uk

Identifying boundaries and repetitions are tasks that are fundamental to music perception and to formal analysis. Trying to replicate listeners' analyses algorithmically is an active research problem, and solutions usually rest on basic assumptions about how listeners analyze music, including how they identify boundaries and repetitions. For instance, it is usually taken for granted that boundaries are found by noticing relatively large changes in a musical parameter. We empirically test how well this assumption accounts for actual musical analyses.

We use the SALAMI [Structural Analysis of Large Amounts of Music Information] corpus of 1,240 structural analyses provided by 9 listeners, all pursuing graduate music degrees. Although this data was not collected as part of a psychological experiment, we demonstrate the value of cautiously repurposing it as such. From the recordings, we compute a set of acoustic features related to timbre, harmony, rhythm, tempo and local key, each for a range of time scales, and derive from these a set of points where the recording is changing the most. We find that boundaries were nearly all indicated by changes in one feature or another, and that coinciding with several simultaneous changes increased a point's likelihood of being a boundary. However, such changes indicated countless other points not deemed boundaries by listeners. Occurring near an acoustic change is thus found to be a *nearly* necessary and grossly insufficient condition for a point to be perceived as a boundary.

The data also revealed differences between how often changes in different musical features at different time scales correctly indicated boundaries, and the genres for which boundaries were more or less predictable. In a final experiment, we compare boundary profiles generated from the novelty functions and from the annotations, and discover that a boundary's salience appears correlated to how acoustically novel it is.