Modeling Music and Memory using Information Retrieval Techniques and Games with a Purpose

Jan Van Balen†, Themistoklis Karavellas, John Ashley Burgoyne*, Frans Wiering†, Dimitrios Bountouridis†, Henkjan Honing*, Remco C. Veltkamp†*

† Dept. of Information and Computing Sciences, Utrecht University, the Netherlands
* Music Cognition Group, University of Amsterdam, the Netherlands

In this abstract, we give an overview of a large interdisciplinary project on musical memory and musical heritage. Two studies are described in which these topics are explored from a computational humanities perspective as well as the common thread that will link these two studies as the project continues.

The types of musical variation occurring in large collections of popular and folk music are an ideal subject for computational musicologists, e.g. for the analysis of oral and cultural transmission and for better understanding the nature of musical similarity. With an interest in applying current methods from folk music analysis to popular music, we have created the Cover-Song Variation (CSV) dataset, an expert-annotated dataset of cover versions of pop melodies (http://www.projects.science.uu.nl/COGITCH/CSV). The dataset consists of a subset of the Second Hand Song dataset, with annotated start times and melody transcriptions for 240 short sections that correspond across different renditions.

The notion of music similarity has been explored to great extent in the field of Music Information Retrieval (MIR), but the computational models used most in MIR often deploy abstract representations of the audio that allow little or no cognitive, perceptual, or otherwise human interpretation, and are thus ignored by researchers outside MIR. For this reason, we have developed two new types of audio features for the assessment of melodic and harmonic similarity that are not only compact and indexable (for MIR) but also interpretable (for heritage and cognition research). The features show reasonable performance in a recent retrieval.experiment.

When it comes to studies of musical memory and popular music, the first concept that comes to mind is the hook. We developed a musical game-with-a-purpose that allows us to locate the parts of a song listeners recognize fastest. In Hooked! (http://hooked.humanities.uva.nl), users go through a three-step process to prove they have recognized a song. We are analyzing reaction times to construct a cognitive model of hooks, in particular to evaluate the relative strength of two leading theories: the one-time-only hypothesis that a single, striking moment is sufficient to imprint a piece of music into long term memory and the competing one-more-time hypothesis that a hook gains power primarily due to repetition.

Ultimately, we aim to test a number of hypotheses about the relation between stability in (oral) transmission and long-term mnemonic salience. Does a model of hooks improve retrieval of folk song variations? Do variations occur less in parts of songs that are better remembered? To answer these questions, we have based the CSV dataset and the Hooked! dataset on a common list of 2000 well-known pop songs in Benelux and have cross-referenced the start times of the structural sections of the common songs in the data sets.