Using Computational Models to Harmonise Melodies

Raymond Whorley, Geraint Wiggins, Christophe Rhodes, and Marcus Pearce*

Centre for Cognition, Computation and Culture
Goldsmiths, University of London
New Cross, London SE14 6NW, UK.
*Wellcome Laboratory of Neurobiology
University College London
London WC1E 6BT, UK.
{r.whorley,g.wiggins,c.rhodes}@gold.ac.uk
marcus.pearce@ucl.ac.uk

Abstract. The problem we are attempting to solve by computational means is this: given a soprano part, add alto, tenor and bass such that the whole is pleasing to the ear. This is not easy, as there are many rules of harmony to be followed, which have arisen out of composers' common practice. Rather than providing the computer with rules, however, we wish to investigate the process of learning such rules. The idea is to write a program which allows the computer to learn for itself how to harmonise in a particular style, by creating a model of harmony from a corpus of existing music in that style. In our view, however, present techniques are not sufficiently well developed for models to generate stylistically convincing harmonisations (or even consistently competent harmony) from both a subjective and an analytical point of view. Bearing this in mind, our research is concerned with the development of representational and modelling techniques employed in the construction of statistical models of four-part harmony. *Multiple viewpoint systems* have been chosen to represent both surface and underlying musical structure, and it is this framework, along with *Prediction by Partial Match* (PPM), which will be developed during this work. Two versions of the framework have so far been implemented in Lisp. The first is the strictest possible application of multiple viewpoints and PPM, which reduces the four musical sequences (or parts) to a single sequence comprising compound symbols. This means that, given a soprano part, the alto, tenor and bass parts are predicted or generated in a single stage. The second version allows the lower three parts to be predicted or generated in more than one stage; for example, the bass can be generated first, followed by the alto and tenor together in a second stage of generation. We shall be describing and demonstrating our software, which uses machine learning techniques to construct statistical models of four-part harmony from a corpus of fifty hymn-tune harmonisations. In particular, we shall demonstrate how these models can be used to harmonise a given melody; that is, to generate alto, tenor and bass parts given the soprano part. Output files are quickly and easily converted into MIDI files by a program written in Java, and some example MIDI files will be played.