Functional Representations of Music

James McDermott*, University College Dublin.

April 30, 2012

Music is an interesting domain for the study of computational creativity. Some generative formalisms for musical composition (e.g. Markov chains) achieve plausible music over short time-scales (a few notes) but appear to be "meandering" over longer time-scales. Imposing a sense of teleology or purpose is an important goal in creating valuable music.

In the field of evolutionary computation (EC), researchers draw inspiration from Darwinian evolution to address computational problems. EC can be applied to aesthetic and creative domains. Although EC is commonly used to generate music, key open issues remain. Formal measurement of the quality of a piece of music in a computational *fitness function* is an obvious obstacle. A naive *representation* for music, such as a list of integer values each corresponding directly to a note, will tend to produce disorganised music.

In previous work, Hoover et al. [1, and later] showed that a *functional* representation could impose organisation and a sense of purpose. In this system, music is represented as a function of time. A fixed piece of pre-existing music is used as a "scaffold": the system then evolves functions, i.e. mappings from the scaffold to new accompanying material. Time-series of numerical "control" variables are also proposed as a means of imposing structure on the music. Fitness is judged interactively.

The XG project is partly inspired by this work. It discards the "scaffold", but relies on the time-series of control variables (see Figure 1).

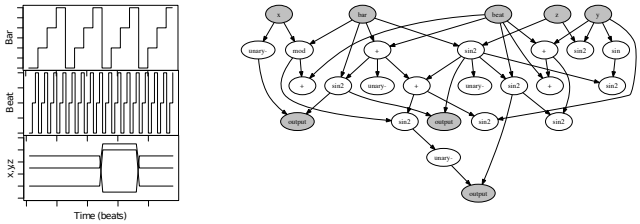

Figure 1: Time-series of control variables (left) impose a bar/beat structure and an overall AABA structure. The evolved function (right) maps these variables to numerical outputs, once per time-step. The outputs are interpreted as music.

It also differs in its internal representation for the mappings (a simple language of arithmetic functions, with special *accumulator* functions at the outputs to control volume), and its use of a computational (non-interactive) fitness function. Surprisingly good results arise using this representation in combination with a simple fitness function which rewards *variety* in the output music. Neither the functional representation nor the fitness function is alone capable of producting good results. More details are available in a full paper [2] and online[1].

A longer-term goal of the XG project is to create large-scale musical works as mappings from pre-existing time series arising in nature and human affairs, and from non-musical artforms such as film or still images with a sequential aspect.

[1] Amy K. Hoover, M. P. Rosario, and Kenneth O. Stanley. Scaffolding for interactively evolving novel drum tracks for existing songs. In *Proceedings of EvoWorkshops*, volume 4974 of *LNCS*, page 412. Springer, 2008.

[2] James McDermott and Una-May O'Reilly. An executable graph representation for evolutionary generative music. In *GECCO '11*, Dublin, 2011.

*Funded by IRCSET/Marie Curie; jamesmichaelmcdermott@gmail.com

[1] http://www.skynet.ie/~jmmcd/xg