Evolutionary Rhythm Composition with Trajectory-based Fitness Evaluation

John Huddleston and Dr. Jianna Zhang

Computer Science Department
Western Washington University, 516 High Street, Bellingham, WA, 98229
huddlej@cc.wwu.edu, Jianna.Zhang@wwu.edu

Evaluating creativity in musical genetic algorithms (GAs) requires a balance between objective and subjective fitness measures. This research investigates the use of user-defined compositional rules for composition of rhythmic pieces of any length. Complete pieces of music can be created according to personal definitions of creativity without direct human involvement in the evolutionary process.

Musical GAs are grouped by their methods of fitness evaluation and the scope of the music they attempt to create. The fitness method can be *objective* (Dostál 2005) (Horowitz 1994) (Papadopoulos and Wiggins 1998), *subjective* (Biles 1994), or *hybrid* (Birchfield 2003) (Jacob 1995) (Todd and Werner 1999) (Tokui and Iba 2000) evaluators. Regardless of the fitness methods used, the majority of musical GAs deal with simplified representations of the search domain by either focusing on generation of rhythmic fragments (Dostál 2005) (Horowitz 1994) (Tokui and Iba 2000) or brief melodic phrases for use in solos or later combination with other phrases (Biles 1994) (Jacob 1995) (Papadopoulos and Wiggins 1998) (Todd and Werner 1999). Few projects have attempted to evolve complete compositional structures with many musical elements (Birchfield 2003). This research is a continuation of work done in the evolution of complete pieces using hybrid fitness evaluation.

While previous research on musical GAs has explored the range of possible fitness measures from the subjective (Biles 1994) to the objective (Papadopoulos and Wiggins 1998), most research shows that a hybrid of these extremes provides the best results (Burton and Vladimirova 1999).

Subjective fitness evaluation depends on some form of human input to assign fitness values to genetic individuals during the evolutionary process. Human evaluation reduces the complexity of the GA by removing the need to algorithmically parameterize the user's musical preferences. Initial results from subjective evaluation frequently match evaluator preferences, but many generations of repetitive analysis tends to lead to skewed measurements from evaluator fatigue. In general, human evaluation is recognized as a major performance bottleneck (Burton and Vladimirova 1999).

Objective fitness measures parameterize subjective quality into objective rules which allow the GA to operate creatively without human input (Burton and Vladimirova 1999). Papadopolous and Wiggins, and Horowitz before them, use collections of rules to limit the search domain to palatable results. The rhythmic elements of these GAs' individuals are evaluated based on length of silences between notes, where notes fall in relation to the downbeat, beat density, and beat repetition (Papadopoulos and Wiggins 1998) (Horowitz 1994). Dostál's GA (2005) depends on an external source of previously composed music for objective evaluation. GAs with objective fitness measures require all objective rules to be statically defined before evaluation can begin. The quality of the resulting individuals depends greatly on which and how many rules are defined. Too few rules will lead to noise while too many rules will result in overly simplistic music. An overabundance of rules can also result in prohibitive computational complexity (Birchfield 2003).

Although a wide variety of hybrids exist, two classes of which are the most prominent: trained and evolutionary evaluators. In the case of trained evaluators, an additional machine learning tool, such as a neural network or another GA, is trained to model the user's preferences either through a preprocessing stage with previously selected music (Jacob 1995) or during the evolutionary process itself (Tokui and Iba 2000). Eventually, the learning tools replace the human evaluator thereby automating the evolutionary process. The human must still provide active feedback for a minimum period of time to get the best results. The creativity allowed by the learned rules always depends in part on the initial training set (Jacob 1995).

Evolutionary evaluators use an initial set of rules which are passed into the GA either as part of the population or as their own population. Birchfield (2003) encodes fitness rules in a multi-level structural representation such that each structural level contains the rules for the level below. The appropriate initial rules defined by the user allow the system to evolve its own fitness rules along with the population of results. Todd and Werner (1999) employ co-evolution of a male and a female population in which the females select males according to the quality of the male songs. As with Birchfield's system, the initial conditions provided for the female selector population determines the eventual success (Todd and Werner 1999).

To isolate the effectiveness of the fitness methods from potential problems caused by an exponentially complex search domain, this research focuses solely on the generation of rhythmic music. An “individual” in a population is a two-dimensional matrix or *pattern* with a

Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

cell, $p[x][y]$, denoting the beat value of instrument x at time y in pattern p as in Tokui and Iba's GA (2000). The matrix row $p[x]$ represents an instrument part or instrument row.

Fitness evaluation is loosely modeled on (Birchfield 2003) with a level of user-defined fitness trajectories for each pattern composing an entire musical piece. The evaluation of each trajectory for a given individual returns a fitness value based on the square of the difference between the expected values defined by the trajectory and the actual values defined by the individual. The overall fitness of the individual is the sum of these trajectory fitness values.

Fitness trajectories are classified as either column or row trajectories. Column trajectories define the expected values for all instruments as a function of time. A sample column trajectory might define the number of instruments playing simultaneously or the average beat volume at any given time. Each column trajectory is defined as a cubic Bézier curve. The desired initial and final values of each trajectory over the pattern are specified as well as two control points which shape the Bézier curve (Figure 1.).

Figure 1. A Screenshot for Setting a Fitness Trajectory

Row trajectories define the expected interactions amongst a set of instruments. A rule for instrument rows 1 and 3 to play the same part would be defined by the "Unison" row trajectory. In addition to defining which rows should be included in the trajectory, a "master row" must be selected to act as the example for the other rows in the trajectory. The master row does not have to be in the trajectory row set, but the other rows must follow its lead.

Three trajectories were tested including one column trajectory and two row trajectories. "Temporal Beat Density", the column trajectory, defines the number of simultaneous beats that should occur at any given time. "Unison", the first row trajectory, defines the rows which should all play the same part. The second row trajectory, "Double Rhythms", defines the rows which should play the same parts circularly shifted to the left or right by any amount. Thus, the valid solutions for "Unison" are a subset of the solutions for "Double Rhythms".

Four sets of tests were run with the three trajectories. Each test set ran 20 unique iterations of 50 generations to generate 8-beat-long patterns of 4 instruments (8x4 patterns). The first test set used only "Temporal Beat Density". The second test set added a "Unison" trajectory to the first test set. The third test set added a "Double Rhythms" trajectory to the first set. The final test set consisted of all three trajectories with "Unison" applied to the first two instrument rows and "Double Rhythms" applied to the last two instrument rows.

The four tests show mostly what one's intuition would expect: the simple beat density trajectory is easily solved by itself, the "Unison" and "Double Rhythms" trajectories are more difficult to solve than the single column trajectory, and all trajectories together are the most difficult. Surprisingly the "Unison" and "Double Rhythms" trajectories are solved equally easily despite the decreased size of the "Unison" search domain. This result is most likely due to the simplicity of the patterns evolved.

We have shown that a user-defined compositional structure can be easily solved by a simple GA in piece-meal to construct complete rhythmic pieces. Future work will investigate the relative weighting of trajectories versus a strict Pareto-optimal comparison and the effects these comparisons have on overall population diversity.

References

Biles, John A. "GenJam: A Genetic Algorithm for Generating Jazz Solos". Available: http://www.it.rit.edu/~jab/GenJam94/Paper.html, 1994.

Birchfield, David. "Generative Model for the Creation of Musical Emotion, Meaning, and Form", *Proceedings of the 2003 ACM SIGMM workshop on Experiential telepresence*, 2003, 99-104.

Burton, A. R. and Vladimirova, T. "Generation of Musical Sequences with Genetic Techniques", *Computer Music Journal 23*(4), 1999, 59–73.

Dostál, Martin. "Genetic Algorithms as a Model of Musical Creativity - On Generating of a Human-like Rhythmic Accompaniment". *Computing and Informatics* 22, 2005, 1001-1020.

Horowitz, Damon. "Generating Rhythms with Genetic Algorithms", *Proceedings of the 1994 International Computer Music Conference*. Aarhus, Denmark: International Computer Music Association, 1994.

Jacob, Bruce L. "Composing with Genetic Algorithms". Available: http://www.ee.umd.edu/~blj/algorithmic_composition/icmc.95.html, 1995.

Papadopoulos, G and Wiggins, G. "A Genetic Algorithm for the Generation of Jazz Melodies." *Proceedings of STeP 98*, Jyväskylä, Finland, 1998.

Todd, P.; Werner, G.M., Frankensteinian methods for evolutionary music composition in Griffith, N. and P.M. Todd, Musical networks : parallel distributed perception and performance. 1999, Cambridge, Mass.: MIT Press. xv, 385.

Tokui, Nao, and Iba, Hitoshi. "Music Composition with Interactive Evolutionary Computation". Third International Conference on Generative Art, 2000.