

Demo: A Computer-Assisted Approach to Composing with MaestroGenesis

Paul A. Szerlip, Amy K. Hoover, and Kenneth O. Stanley
Department of Electrical Engineering and Computer Science
University of Central Florida
Orlando, FL 32816-2362 USA
{paul.szerlip@gmail.com, ahoover@eecs.ucf.edu, kstanley@eecs.ucf.edu}

Abstract

This demonstration presents MaestroGenesis, a program that helps users create complete polyphonic musical pieces from as little as a simple, human composed monophonic melody. MaestroGenesis creates music by exploiting two key ideas behind the functional scaffolding for musical composition (FSMC) approach: (1) that music a function of time and (2) that functional transformations of initial human starting melodies, or *scaffolds*, inherit some of the essential human qualities contained in the scaffold. Music in FSMC is represented as a functional relationship between the scaffold and a generated accompaniment. The GUI helps users evolve these functions by importing and developing their music through a breeding process akin to animal breeding, called interactive evolutionary computation. Some resulting pieces are indistinguishable from completely human-composed pieces.

MaestroGenesis Overview

MaestroGenesis is a program that helps users create complete polyphonic pieces with only the musical expertise necessary to compose a simple, monophonic melody. Building on NEAT Drummer, a program that generated drum patterns for existing compositions, MaestroGenesis users begin creating accompaniments by establishing the starting melody, called a *scaffold*, that will provide the initial rhythmic and harmonic seed for the accompaniment (Hoover and Stanley 2009). This scaffold can be any monophonic or polyphonic piece composed by the user in MIDI form. The accompaniment is then represented as a functional transformation of this original scaffold through a method called functional scaffolding for musical composition (FSMC) (Hoover et al. 2012). FSMC exploits the structure already present in the human-composed scaffold by computing a *function* that transforms its structure into the accompaniment.

Because there can be many appealing transformations of any given scaffold, MaestroGenesis encourages users to develop accompaniments through a process similar to animal

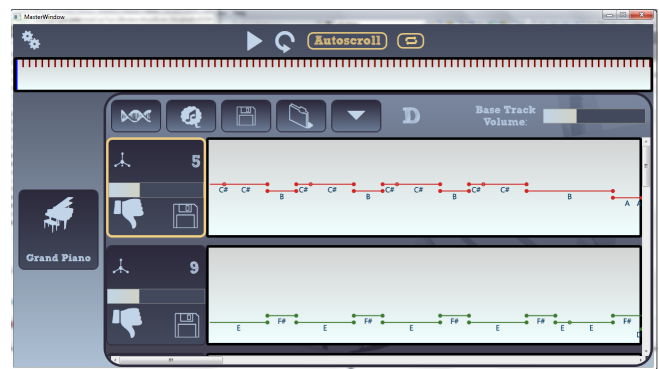


Figure 1: **MaestroGenesis Candidate Accompaniments.** Accompaniments in MaestroGenesis are evolved through a process similar to animal breeding. Candidate accompaniments are evolved by choosing favorites among a population of ten candidates in an interactive process in which each subsequent generation inherits traits from the previous generation.

breeding called interactive evolutionary computation (Takagi 2001). Once the scaffold is chosen, a population of ten accompaniments is displayed. Each is rated as good or bad by pressing the “thumbs-up” button (figure 1). By rating favorable accompaniments higher than less appealing accompaniments, the next generation of accompaniments tends to possess similar qualities to the well-liked parents. Through interactively evolving these accompaniments, they increasingly reflect the personal inclinations of the user.

Unlike previous approaches in assisted music composition (Keller et al. 2006; Simon, Morris, and Basu 2008), an interesting aspect of MaestroGenesis and the underlying FSMC method is that the representation itself requires no explicit musical knowledge. In fact, the only musical constraints are that the output stays within the user-defined key. In the accompaniments shown, the smallest note unit is restricted to an eighth note, but there is no max-




Figure 2: **Visual Representation of Music.** The visual representation of music in MaestroGenesis is different than standard musical notation and is therefore easier for amateur musicians, who may have difficulty reading music. Each horizontal line represents a new note while their varying heights describe a relative pitch.

imum note length. Furthermore, the transformations are evolved with the NeuroEvolution of Augmenting Topologies (NEAT) method for evolving neural networks, in which the genetic mutation and crossover operators are domain independent (Stanley and Miikkulainen 2002).

Selecting Accompaniments

Listening to generated accompaniments is critical for selecting appealing pieces, but aurally assessing each candidate accompaniment requires a significant time investment. Corresponding visual representations, however, can significantly decrease evaluation time. For instance, if a user is looking for a walking bass line and a quick visual inspection reveals that the piece is stationary, i.e. on the same note throughout the piece, the user can quickly discard the candidate. MaestroGenesis allows users to both see and hear accompaniments for quality assessment.

Rather than representing the pieces in standard musical notation, which can impede amateur musicians, accompaniments in MaestroGenesis are viewed as a series of straight horizontal lines over time as shown in figure 2. The level of these lines rises and falls as pitches rise and fall in the accompaniment and the lengths of the lines indicate duration. For those with more musical expertise, the letter name of each note is also displayed. If the accompaniment is played, the current note is highlighted as the progression scrolls in real time.

Timbre selection can play a major role in perceived accompaniment quality. MaestroGenesis addresses this issue by allowing users to adjust timbres in real time with the interface shown in figure 3. Users select from among 16 instrument categories (e.g. bass, guitar, piano) and 128 instruments.

Applications of MaestroGenesis

Because each generation of accompaniments in MaestroGenesis depends on the user's tastes and preferences, accompaniments vary depending on the user and his or her approach to composition with MaestroGenesis. For example, counter-melodies, such as the one in figure 4, can be easily created. This extra harpsichord accompaniment was created for the already complete song, Nancy Whiskey, which was originally composed by Barry Taylor and used with his permission. Because of the information in the three different instruments in the scaffold, MaestroGenesis was able to create




Figure 3: **Instrument Selection.** Accompaniments in MaestroGenesis are first generated in MIDI and then mapped to MP3 sound fonts, thereby allowing the user to quickly alter the timbre qualities of any generated accompaniment. Timbres in MaestroGenesis are organized by instrument category (e.g. bass, guitar, piano, etc.) while each category contains several instances of the category.



Figure 4: **Counter Melodies with MaestroGenesis.** Because MaestroGenesis does not contain explicit musical knowledge, it can easily generate original accompaniment for different styles. The second line on the staff above contains its own melody even though it was evolved based on the melodic line above it. The first line comes from Nancy Whiskey, which was originally sequenced by Barry Taylor and used with his permission. The piece can be heard at <http://eplex.cs.ucf.edu/fsmc/mume2012>.

simple functional transformations that resulted in accompaniments with high perceived quality (Hoover, Szerlip, and Stanley 2011a).

However, MaestroGenesis is also beneficial for amateur musicians who may only possess the ability to compose a simple monophonic melody on their own. Figure 5 shows six instrumental parts, five created by MaestroGenesis from the simple, human-created melody displayed as Layer 1. For each line of accompaniment, users are permitted to change the underlying scaffold, which has varying effects on the compositions. For example, Layer 2 comes from evolution based only on Layer 1, and has clear rhythmic and pitch similarities. Layers 3, 4, and 5 are based on runs with Layer 1 and Layer 2 as the scaffold. However, Layer 5 is from a separate run and more resembles the starting melody.

More results from the MaestroGenesis system are avail-

Figure 5: **Creating Polyphonic Accompaniments.** Users can generate any number of instrumental parts to accompany a single monophonic instrument. The melody in Layer 1 was composed by Marie E. Norton, an undergraduate independent study student and member of the MaestroGenesis team. That melody scaffolded the remaining instrumental parts of the piece that she later evolved with MaestroGenesis. The piece, which is displayed in concert pitch (i.e. all of the parts are written in the same key), can be heard at <http://eplex.cs.ucf.edu/fsmc/mume2012>

able in Hoover et al.; Hoover, Szerlip, and Stanley; Hoover, Szerlip, and Stanley. These papers present more user-generated pieces and include studies that explore musical quality.

Conclusion

MaestroGenesis helps users compose polyphonic pieces from simple starting melodies through functional scaffolding for musical composition (FSMC) (Hoover et al. 2012). Users begin with the starting melody and interactively evolve musical accompaniments in a process similar to animal breeding. While creating accompaniments, users can change the starting melody for more variety or they can even adjust the accompaniment timbre in real time. The program is available for download at <http://maestrogenesis.org>.

Acknowledgments

This work was supported in part by the National Science Foundation under grant no. IIS-1002507 and also by a NSF Graduate Research Fellowship. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

References

Hoover, A. K., and Stanley, K. O. 2009. Exploiting functional relationships in musical composition. *Connection Science Special Issue on Music, Brain, & Cognition* 21(2):227–251.

Hoover, A. K.; Szerlip, P. A.; Norton, M. E.; Brindle, T. A.; Merritt, Z.; and Stanley, K. O. 2012. Generating a complete multipart musical composition from a single monophonic melody with functional scaffolding. In *Proc. of the 3rd Intl. Conf. on Computational Creativity (ICCC-2012)*.

Hoover, A. K.; Szerlip, P. A.; and Stanley, K. O. 2011a. Generating musical accompaniment through functional scaffolding. In *Proceedings of the Eighth Sound and Music Computing Conference (SMC 2011)*.

Hoover, A. K.; Szerlip, P. A.; and Stanley, K. O. 2011b. Interactively evolving harmonies through functional scaffolding. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2011)*. New York, NY: The Association for Computing Machinery.

Keller, R. M.; Morrison, D.; Jones, S.; Thom, B.; and Wolin, A. 2006. A computational framework for enhancing jazz creativity. In *Proceedings of the Third Workshop on Computational Creativity, ECAI 2006*.

Simon, I.; Morris, D.; and Basu, S. 2008. Mysong: Automatic accompaniment generation for vocal melodies. In *Proc. of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems*, 725–734. ACM.

Stanley, K. O., and Miikkulainen, R. 2002. Evolving neural networks through augmenting topologies. *Evolutionary Computation* 10:99–127.

Takagi, H. 2001. Interactive evolutionary computation: fusion of the capabilities of ec optimization and human evaluation. *Proceedings of the IEEE* 89(9):1275–1296.