

Harnesses artificial life to evolve original music

Composer scores advance in high-tech tunes

By R. Colin Johnson

PARIS – Just as IBM's Deep Blue showed the world a computer can play chess as well as a human master, Eduardo Reck Miranda, a researcher for the Sony Computer Science Laboratory here, aims to demonstrate a computer program able to compose original music. So far, neural networks have succeeded in imitating distinct musical styles, but truly original compositions have remained elusive. Miranda is tackling that problem with an orchestra of virtual musicians—agents—who interact to compose truly original music.

“From the viewpoint of a composer, I can hardly say that our agents are composing music at this stage,” Miranda acknowledged. Instead, he said, “The breakthrough in this work is the action of collective machine learning for generative music systems.” Rather than modeling specific musical styles or compositional processes, the new approach is a first step toward collective music-making with virtual musicians, each of which learns how to compose and play individually.

The Brazilian-born Miranda composes chamber and electro-acoustic pieces, compositions that have won prizes in the Americas, Asia and Europe. Today, at Sony, he specializes in artificial intelligence in music. In his latest book, *Composing Music with Computers* (Focal Press), Miranda summarizes his AI research, which began with cellular automata and evolved into an “adaptive games” strategy based on artificial-life models.

“Perhaps one of the greatest achievements of artificial intelligence today lies in the construction of machines that can compose music of incredibly high quality,” he said.

But AI's achievement is restricted to mimicking the style of existing composers, either with a set of AI rules, or by learning a composer's style with a neural network. In other words, computers can compose a new Bach cantata, but they will never com-

pose anything novel, because their algorithms merely encapsulate a particular style of music.

For a computer to create truly novel compositions, Miranda has turned to artificial life (AL) models—the new fodder for what he calls evolutionary musicology. The computer becomes endowed with an adaptive set of music-composition tools and processes. Under the control of an artificial life model the set displays “emergent” behaviors—that is, truly novel compositions. Miranda has explored this approach with cellular automata, genetic algorithm-like processes he summarized as “adaptive games.”

He began his investigation with musical pieces composed using cellular automata, and created two computer programs to assist in the process—a software synthesizer called ChaoSynth and a generator of musical passages called Camus.


Cellular automata are organized as a set of identical “cells” in a two-dimensional grid. Each has a set of rules to follow, but can only communicate with its nearest neighbors. The overall strategy is to begin with a simple set of rules affecting only local changes, but to evolve an overall organization that exhibits the emergence of new high-level structures—in this case, music.

The process begins with a random distribution of cell states, then evolves a pattern through repeated application of each cell's rule set in each time period. After an epoch of time periods, a pattern spontaneously emerges that is plain to see with the eye, as in the Game of Life, but is translated into a soundscape by Miranda's ChaoSynth.

ChaoSynth uses a granular synthesis method to translate the two-dimensional “Game of Life” pattern into a signal that can drive a speaker. By accumulating a sequence of short sound “grains”—typically 10 to 100 microseconds in length—many different cells can contribute to the final sound sample. Granular synthesis is often used to permit the pitch and duration of a sound


to be adjusted independently, but the pitch and amplitude of Miranda's sound grains were set by the rules being executed by each cell at any given time.

“ChaoSynth is a successful system because it can synthesize a large amount of unusual sounds that are not normally found in



stead of noise. For instance, he decided imitation would be the primary directive of his agents, because among animal life, imitation is counted as the first steps toward “intelligence.”

“Life is everywhere animated with mutual interests between individuals who interact with



Cellular automata in musical composition starts with a random cell pattern, left, and evolves into the more-articulated form, right, using a specific set of rules.

the real acoustic world, but which nonetheless sound pleasing to the ear,” said Miranda.

The ChaoSynth “soundscapes,” often accompanied by Miranda on the piano, have won awards for being “musical” despite the unearthly sound of their individual parts. While pleasing to the ear when they are blended in composition, those everyday sounds fit no known category when isolated.

Musical evolution

Miranda concluded, after years of research building generative music systems with cellular automata, that he needed to extend his theory to account for the social formation of music and its cultural evolution. Accordingly, current musical research focuses on interacting communities of distributed agents, each with motor, auditory and cognitive skills. The goal is to evolve a repertoire of melodic patterns in the memories of the agents, by virtue of spontaneous creation, adaptation and reinforcement.

So far, the results have not been as “musical” as his cellular automata compositions, but he is being patient because the motivation for his current work is broader. Miranda has had to conjure the innate capabilities that make music sound like music in-

one another to express reciprocity of intentions and for coordinating tasks,” Miranda said. “Up to now, I considered this work more at a theoretical level, as an investigation into the origins of music, rather than as a practical method to actually compose music. Only now have I started composing using the sounds created by adaptive agents.”

To begin, Miranda equipped his agents with a voice synthesizer, hearing apparatus, memory device and an “enacting script” that the system follows to communicate with other agents. The agents compute the parameters of their synthesizer—principally pitch and duration—and play them to each other. The agent's memory stores its “song” repertoire and can extract the pitch of songs it hears to compare them with its own repertoire.

A variety of statistical, threshold and reinforcement parameters enables the agents to adapt their behaviors toward the goal of imitation. For instance, just because an agent “hears” a song by tracking its pitch, it cannot exactly imitate the sounds it hears because it is saddled with its own motor-skill limitations. In fact, each agent's motor skills also adapt, so that over time agents become more skillful at imitating songs similar to

ones already in its repertoire.

To embody this behavior, agents have many different “memories.” A perceptual representation memory, for instance, holds the pitch information it “hears” while a separate motor-map memory describes the skillfulness of its synthesis abilities. Ironically, despite being able to imitate what they have already learned, they are nevertheless driven to crave novelty.

In practice, the agents pair off so that one can “listen” while the other is “playing” one of the songs from its repertoire. The listening agent then compares the song it perceives with those already in its memory and plays the closest match back to the original agent. Then the roles reverse, with the agent that originally played the first song, comparing its perception of the song it got in return with the original it played. If it is a close match, then the originating agent will play it back again unaltered, and the imitating agent will reinforce its learning of the song. However, if a different song is a closer match, then the original agent will remain silent, and the imitating agent will discard that association rather than reinforce it.

Agents keep track of the “popularity” of each of their songs by tracking how many other agents have successfully imitated them. Songs rejected by one imitator but accepted by many other imitators are left alone. But songs that are repeatedly “rejected” by imitators are eventually deleted and replaced with a new song that is a permutation of the rejected song.

For Miranda, the novelty here is that his adaptive games can learn to actively compose music, not just passively generalize existing musical knowledge. If borne out, this hypothesis will lead to realistic musical cultures that can be evolved by furnishing virtual musicians with what he calls “the proper cognitive and physical abilities, combined with appropriate interaction dynamics and adequate environmental conditions.”