Virtual Music (review)

Michael Theodore

Computer Music Journal, Volume 26, Number 4, Winter 2002, pp. 92-95 (Review)

Published by The MIT Press

pitch organization (James M. Baker, Providence, Rhode Island), on musical imagery during composition processes (Rosemary Mountain, Montreal, Canada), and on musical imagery in performance, composition, philosophy, and cosmology in Indian music (Lewis Rowell, Bloomington, Indiana).

One may claim that some of the articles do not go deep enough into specific research on musical imagery, or that some authors have used only a few subjects for their empirical investigations. However, we need to keep in mind that this is the most comprehensive publication on this topic to date, and that all research is well-documented and referenced so that readers can locate additional literature easily. (Name and subject indices at the end of the volume help immensely in that respect.) The strength of the book in general is, as mentioned above, its thorough introduction to the manifold problems and questions related to musical imagery. The book seems most useful not just for musicologists and music theorists, but also—or better, especially—by music educators, composers, and performers.

David Cope: Virtual Music

Hardcover, 2001, ISBN 026203283X, 292 pages, illustrated, appendices, audio CD, US$ 45.00; The MIT Press, Five Cambridge Center, Cambridge, Massachusetts 02142-1493, USA; telephone (+1) 800-356-0343; electronic mail mitpress-order @mit.edu; Web mitpress.mit.edu/.

Reviewed by Michael Theodore Boulder, Colorado, USA

David Cope's newest book, *Virtual Music*, arose out of a weekend of papers, panels, concerts, and discussions devoted to Mr. Cope's extraordinary Experiments in Musical Intelligence (EMI) software. EMI is a program which, when given a set of compositions by a particular composer (or composers) as input, attempts to autonomously compose new pieces in the style of the source music. The weekend's events included a highly distinguished panel of presenters, including Mr. Cope, Douglas Hofstadter, Eleanor Selfridge-Field, Bernard Greenberg, Steve Larson, Jonathan Berger, and Daniel Dennett, each of whom also contributed at least one chapter to the book.

Virtual Music is divided into four parts. The first provides a context for EMI by giving a historical overview of algorithmic composition, and also includes an informal description of the mechanics of EMI (given by Mr. Hofstadter). The second part provides a detailed "case study," demonstrating the composition of an EMI work from beginning to end. The third part consists of scholarly evaluation and commentary on the program. The concluding section includes Mr. Cope's response to the criticisms offered by the other scholars, as well as his speculations on the directions the software might take in the future. The book's multiple appendices contain generous amounts of musical examples, and an audio CD of EMI compositions is included as well.

Mr. Hofstadter's segment in the first part of the book is one of the highlights, both for his excellent overview of EMI and for the humorous manner in which he is able to raise some of the troubling philosophical questions that the software provokes. Mr. Hofstadter reduces the essential operation of the program (when given a set of pieces to operate on) to two actions: 1) chop-up, then 2) reassemble. In the "chop-up" phase, the program dynamically segments the input music into meaningful units (on several levels of hierarchical structure). This is no easy task, as the musical material must be chopped up finely enough so that the end result doesn't overly resemble the original sources, but not so fine that the musical coherence becomes lost. In the "reassemble" phase, the program constructs a piece of music by recombining the fragments, attempting both to create a coherent flow on the local level and to ensure that the global patterning of fragments resembles that of the source music.

A number of underlying principles guide the program through these two main stages. The local flow is achieved in part through strong voice-leading rules. As Mr. Hofstadter describes, "Imagine that we have just inserted a fragment *f1*, and are considering whether to insert fragment *f2* right after it, drawn from somewhere [else] in the input." EMI's voice-leading rules stipulate that "the initial note of the melodic line of fragment *f2* should coincide with the next melodic note to which fragment *f1* led in the original context. In other words, a given fragment's melodic line should link up smoothly with the melodic line of its successor fragment."

The succession of fragments is also guided by a framework of "tension-resolution," which EMI quantifies by attaching one of the letters S, P, E, A, or C to the fragment. The letters stand for Statement, Preparation, Extension, Antecedent, and Consequent. This framework attempts to capture where on the tension-resolution continuum the fragment is situated. EMI determines the appropriate label for a given fragment by examining such things as the level of dissonance in the sonority as well as the metrical

placement of the fragment. The software attempts to determine the tension-resolution status of a fragment not only on the local level but also on multiple hierarchical levels (the level of the phrase, of the period, of the section, etc.). Mr. Hofstadter sums up the core local and global processes of EMI as follows: 1) ''Sequential assembly of fragments that have the highest possible degree of agreement of SPEAC labels on all hierarchical levels,'' and 2) ''Stitching-together of fragments so as to respect voice-hooking constraints and so as to match local textures.''

EMI also takes steps to mitigate against the possibility that the segmentation process might have disassembled important musical patterns that extend beyond the boundaries of the resultant sections. One of the most important of these types of patterns is the ''signature,'' which Mr. Cope defines as ''contiguous note patterns which recur in two or more works of a single composer and therefore indicate aspects of that composer's musical style. Signatures are typically two to five beats (four to ten melodic notes) in length and usually consist of composites of melody, harmony, and rhythm. Signatures typically occur between four and ten times in any given work.'' Signatures can be thought of (perhaps crassly) as special ''licks'' that are dear to a particular composer. The inclusion of signatures immediately causes EMI's output to sound considerably more convincing (this is especially true in the case of ''signature-happy'' composers such as Wolfgang Mozart).

EMI has several other higher-level principles guiding its process, including ''earmarks,'' which are patterns that announce upcoming important structural events (such as the cadential trills in Mozart's piano concerti that prepare listeners for the upcoming cadenza), and ''unifications,'' which are patterns of local importance that are infused throughout the work being composed. All of these various constraint processes operate in parallel, with much cross-talk between the various levels of structure.

After roughing in the mechanics of the program, Mr. Hofstadter asks what is perhaps the most important question: how well does the program perform? More precisely, does the program create new works that are convincingly in the style of the works in its database? As the final product of the program is an aesthetic artifact (music), this is necessarily a subjective question, and it becomes even harder to answer on account of the biases that people invariably bring to the experience of listening to computer-composed music. People listen with a completely different set of ears when told beforehand that the music has been composed by computer (by listening for the types of awkward passages they think a computer will create, by assuming that there won't be a perceptible ''emotional'' core, etc.). Therefore, when giving public presentations, both Mr. Cope and Mr. Hofstadter typically present the musical output of the software in the form of the ''The Game,'' an example of which follows. The audience is told that they will be hearing several mazurkas in the style of Frederic Chopin (for instance). They are also told that at least one of the pieces was composed by EMI, and that at least one of the pieces was composed by Chopin. After hearing the compositions, the audience members vote either ''Chopin'' or ''EMI'' for each piece (disqualifying themselves if they happen to already know that a particular piece is by Chopin). The results are telling. When EMI is at its best, most audiences, including audiences comprised primarily of specialists (one account in the book describes a session that took place at the Eastman School of Music, with both students and faculty present), are highly uncertain as to which music was composed by which entity.

''The Game'' thus establishes that something very real is going on here. Mr. Hofstadter then asks, ''what does this mean?'', and finds himself unsettled by the implications. He had previously expressed great reservations regarding the composition of music by machines, as in the following passage from his landmark book, *Gödel, Escher, Bach*:

> Question: Will a computer program ever write beautiful music? Speculation: Yes, but not soon. Music is a language of emotions, and until programs have emotions as complex as ours, there is no way a program will write anything beautiful. There can be ''forgeries''—shallow imitations of the syntax of earlier music—but despite what one might think at first, there is much more to musical expression than can be captured in syntactical rules.

Hofstadter has always had a special affinity for the music of Chopin, and indeed has felt that the experience of listening to this music was akin to receiving messages of the greatest profundity directly from the composer's soul. But the fact that EMI can compose similarly ''emotionally-charged'' music without having a soul (or any human attributes) directly challenges Mr. Hofstadter's deeply held view of music. What if music is ''no deeper'' than the manipulation of patterns? What if the ''meaning'' that music seems saturated with is simply an illusion (i.e., what if music simply sounds as if it means something, but really doesn't)? Would this mean that Chopin was in fact not an artist,

but rather a skilled artisan? Mr. Hofstadter spells out these fears in many strains of doggerel.

The contributions of the other scholars as they weigh in on these (and other) issues make for fascinating reading. Mr. Cope himself seems to think that most of Mr. Hofstadter's concerns aren't really an issue. Indeed, he is as uncomfortable talking about things like ''emotional substrate'' as Mr. Hofstadter is obsessed with them. He is also uncomfortable with the idea of ''communication'' in music, and says that he never thinks about ''communicating'' when composing his own (non-EMI) works. Rather, he is interested in ''creating well-balanced structures within which [he] hopes to weave inventive musical ideas.'' Mr. Cope marshals as support for his stance a hyper-formalist quote from Igor Stravinsky who claimed that music is incapable of expressing ''anything at all'' (adding ''expression'' to the list of musical issues that Mr. Cope thinks are smokescreens). Yet, Mr. Cope hastens to add, the inability to express anything doesn't make music ''meaningless.'' This reader wishes that these particular ideas were fleshed out more, for how can music have a ''meaning'' without any ''expression'' of that meaning? If music is incapable of expressing anything at all, have countless composers simply been deluding themselves and their performers by including expressive indications as part of their scores? (Perhaps performers should take an eraser to all of those ''espressivo'' markings!) On a related front, Mr. Cope writes that he ''does not believe that any work of art is intrinsically better than any other work of art,'' and casts aspersions on the ''Western tradition of ascribing 'greatness' to some composers, while other are of lesser quality.'' His statements seem to contradict the very idea of ''quality'' (or at least any sense of quality that is not entirely subjective). Yet clearly some notion of quality drives Mr. Cope's own work with EMI (as he judges some works better than others) as well as his choice of music for the database—after all, he seems to favor ''major'' composers like Bach and Mozart over ''minor'' composers such as Albinoni.

Similarly thorny issues come up in the commentary by Bernie Greenberg, who limits himself to EMI's J. S. Bach emulation. Mr. Greenberg, a believer in ''strong'' Artifical Intelligence (AI), is convinced that ''beautiful, arbitrarily interesting, emotionally challenging music can be created programmatically.'' However, he feels that EMI's Bach-style music is lacking on two interrelated fronts: ''high level emotional architecture,'' and ''low-level contrapuntal technique.'' He cites as an example of both Bach's sublime handling of form in the C Minor Passacaglia, BWV 582.1. The dramatic architecture of the piece includes a steady ratcheting of dramatic intensity, with perfectly placed slackenings along the way. For Mr. Greenberg, this work demonstrates a level of mastery that EMI has not yet attained.

Although Mr. Cope has clearly made great progress in modeling formal architecture, this type of modeling does not address what Mr. Greenberg claims is the distinct notion of ''emotional architecture.'' Mr. Greenberg cites the work of Roger Schank and Robert Abelson on the modeling of drama as an example of the type of work that he believes will lead to the solution of this problem, and suggests the use of a ''state'' network to model this kind of dramatic rhetoric. The network would control, for instance, the degree of harmonic and contrapuntal ''liberty.'' Bach is continually regulating these features of his compositions (they often track the dramatic profile of the work), but EMI does not seem to have this capacity built in. Indeed, Mr. Greenberg finds EMI's Bach-style counterpoint far too ''timid.'' He thinks Mr. Cope's SPEAC model is a good beginning for modeling this specific aspect of music, but that this problem is highly non-trivial and needs much future research.

Jonathan Berger also suggests other important avenues for future research. In particular, he calls attention to the critical role that listeners play in the shaping of a musical experience, and suggests that a connectionist approach (involving neural networks) is the best way to model the cognitive processes that are activated in the listening process. He therefore built (with Dan Gang) a neural network that is intended to be EMI's ''sister''—Experiments in Musical Listening. Mr. Berger's chapter (titled ''Who Cares If It Listens?'') describes the manner in which the network yields insight into the interplay between musical expectation and realization.

Although the authors discussed so far seem to believe that the shortcomings of EMI can be solved with time, Daniel Dennett's intriguing commentary suggests that there are certain problems that may be fatal. His chapter highlights the ''continuity between all sorts of creativity,'' which he believes is ultimately algorithmically based. Life itself was created through a variety of algorithmic processes, and these same processes in turn gave rise to the relatively recent branching between the plants and the animals. He argues that the compositions of EMI are special cases of the same processes that created the compositions of Bach, ''the apples and spider webs, and the organisms that made them.''

Mr. Dennett asserts that all invention and creativity is derived from generate-and-test algorithms. Various combinations are spun out, and then particular elements are chosen out of the many possible instances. In addition, all invention is firmly built upon previous invention. The ''invention'' of homo sapiens (three to four billion years), the invention of human culture (three million years), and 42 years of living were all in place before Bach composed his *St. Matthew's Passion*. Therefore, ''no human being, no matter how great a genius, does all of the creative work that goes into a work of art.''

Mr. Dennett uses the Hofstader-coined term ''spontaneous intrusion'' to describe what he believes is another central element of the creative process. ''In the real world, almost everything that happens leaves a wake, makes shadows, has an aroma, makes noise, and this provides a bounty of opportunities for spontaneous intrusions. It is also precisely what is in short supply in a virtual world.'' The world inhabited by EMI is many orders of magnitude simpler than the world of human musical composition. EMI demonstrates to an astonishing degree just how much can be accomplished in such a ''clean'' environment, but is also perhaps limited by the clarity of the model. Thus, one strategy to more closely model creativity would be to add noise to every component of the program, which would provide the opportunity for the serendipitous creative transformation of noise into signal. However, one eventually reaches a point of diminishing returns, ''for in order to get closer and closer to the creativity of a human composer, your model has to become ever more concrete: it has to model more and more of the incidental collisions that impinge on an embodied composer.'' Mr. Dennet implies that this issue might be a barrier that could prevent EMI (or its offspring) from crossing the threshold into the highest level of musical creativity.

Also included in the book is a humorous letter from Steve Larson to ''Emmy,'' treating the program as if it were a student who should perhaps show up at some of Professor Larson's office hours (and expertly pointing out specific instances in which some of Emmy's Bach works fall short of the ideal), and a thought-provoking chapter by Eleanor Selfridge-Field highlighting the continuity between EMI and earlier historical examples of algorithmic composition.

If you have any interest at all in this field, you should absolutely read this book. It is also a good place to start if you are interested in what you've heard about (or of) Mr. Cope's work, but haven't yet read any of his work. EMI is without question a seminal achievement, and it is fitting that the collection of brilliant thinkers represented in this book came together around its axis. It's true that EMI still has shortcomings, but it really represents just the start of an exciting human/machine journey. How close will we be to ''Chopin's Fifth Ballade'' in 20, 80, or 200 years? No one can really say, but reading *Virtual Music* offers a fascinating glimpse into some of the technical and philosophical questions that will frame the upcoming adventure.

Recordings

Bill Alves: The Terrain of Possibilities

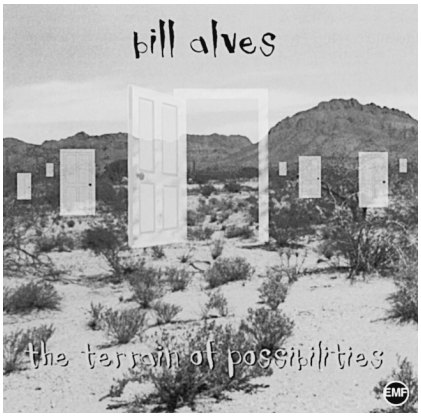

Compact disc, EMF CD 002, 1997; available from Electronic Music Foundation, 116 North Lake Avenue, Albany, New York 12206, USA; telephone (888) 749-9998; (518) 434-4110; fax (518) 434-0308; electronic mail emf@emf.org; Web www.cdemusic.org.

Reviewed by Laurie Radford
Edmonton, Alberta, USA

The Terrain of Possibilities brings together six computer music works from the mid to late 1980s by composer Bill Alves, a professor at Harvey Mudd College in Claremont, California. Mr. Alves has been simultaneously investigating the seemingly unrelated worlds of indigenous world musics and music with computers since that time, bringing them together in a series of works involving computer and instrumental resources as well as specially designed tuning systems. All of the pieces on this compilation were created using a Synclavier II computer music system. The disc offers an example of a mode of making music that served as a stepping-stone to our present-day digital audio workstation-based composition environments. The Synclavier II offered a compositional world in a package, including an attractive hardware performance interface, sound and sequence editing features, and a host of