About This Issue

With the exception of its first article, this issue continues the Summer 2004 issue's theme of music information retrieval (MIR). Three of the articles represent the remainder of the group of eight MIR articles guest-edited by Stephen Bainbridge and Holger Hoos, the other five having been published in the previous issue. These three articles were postponed until this issue owing to limitations of space (i.e., page count) or time (publication deadlines). Another article in this issue, by David Temperley, happens to fit well with the MIR subject matter, but it was actually submitted independently to the Journal.

The issue commences with David Cope's article, which concerns automatic composition and machine learning. Mr. Cope explains how his free software Gradus learns rules of two-voice, note-against-note counterpoint from example musical passages provided as input. Notions of acceptable intervals and voice-leading are inferred, not hard-wired, meaning that the program can emulate non-traditional counterpoint. After deriving rules from the examples, the program composes a new counterpoint against a *cantus firmus* supplied by the user. It back-tracks as necessary to avoid contrapuntal dead ends, but meanwhile it develops rules to eliminate back-tracking, so that the more experience it gains with different *cantus firmi,* the more quickly it composes. According to the author, it is not difficult to adapt his Lisp code to allow complex types of counterpoint. He has in fact created such extensions, showing us a four-voice fugue in florid counterpoint as evidence.

The next three articles each deal with a different category of music analysis: metrical, harmonic, and motivic. David Temperley's article echoes a theme from the past issue: the need for methods of evaluating and comparing different MIR models—in this case, algorithms for metrical analysis (or "beat-tracking"). He lists some two dozen metrical models, including the one he developed with Daniel Sleator (see *Computer Music Journal* 23:1, Spring 1999). The article's mission is not to evaluate all these models, but rather to formulate a technique for conducting such evaluations. The author's proposal encompasses a representation for the input data, a corpus of music for analysis, a set of manually created correct analyses ("goldfiles"), and a means of matching a metrical model's output to the goldfiles in order to arrive at a score measuring the model's performance. The proposed technique was implemented and tested on the author's own model. An evaluation system like this can be put to good service not only as a means of comparing the performance of finished models, but also as a test bed during their development.

Christopher Raphael and Joshua Stoddard present a statistical-learning approach to harmonic analysis. Such analysis has applications in MIR for the knowledgeable user who can construct queries involving harmonic functions or chord progressions. The authors use automatically trained Hidden Markov Models, a broadly useful machine-learning technique, to recognize harmonic functions within the raw pitch and rhythm information obtained from MIDI files. The experimental results show good performance of the authors' program on a small collection of classical music.

Motivic analysis is the goal of the work presented by Olivier Lartillot. He presents a framework for automatic pattern discovery in melodic lines, based on an incremental and adaptive search for repeated patterns of notes along multiple parameters. These parameters include the pitch interval between two successive notes (expressed in semitones), inter-onset ratio (i.e., the relative difference between successive notes' attack times), and pitch contour, a more general characterization of melodic shape. The current implementation, written in Lisp as an OpenMusic library, focuses on the first two of these parameters. Mr. Lartillot seeks to surpass the approach to grouping proposed by Fred Lerdahl and Ray Jackendoff's book *A Generative Theory of Tonal Music* (MIT Press, 1983). The latter relies on strict segmentation and a single

Front cover: An image-processed excerpt from a fugue composed by David Cope's automatic counterpoint software.

Back cover: The front cover of the compact disc *HPSCHD*, which is reviewed in this issue. This recent studio recording of the 1967–1969 multimedia composition by John Cage and computer-music pioneer Lejaren Hiller features a realization of the computer part by Joel Chadabe and harpsichord performances by Robert Conant. The Electronic Music Foundation released the disc as item 138-2 on the EMF label (see www.cdemusic.org).

hierarchy for a given piece of music, whereas Mr. Lartillot's approach allows motives to overlap and to be discovered independently of each other, potentially resulting in a multiplicity of interpretations for the same set of notes. To prevent an excessive number of possibilities, techniques to filter redundancy are employed.

The final article, by Wei-Ho Tsai et al., involves a quite different area of MIR research. Here, the problem domain is how to automatically cluster audio recordings by singer, based on characteristics of the singer's voice. The authors adopt an unsupervised machine-learning approach and present experimental results that show the effectiveness of their system on a set of several hundred tracks of Mandarin popular music.

The reviews in this issue evaluate a major music-industry trade show, a nontechnical book for sound recordists, some synthesis software, and a number of CDs. Composer Larry Austin makes two appearances in this section, first as a reviewer and then, on the other side of the coin, as the subject of a review. Announcements of recent products round out the issue.