Harmonic Navigator: An Innovative, Gesture-Driven User Interface for Exploring Harmonic Spaces in Musical Corpora

David Johnson[1], Bill Manaris[1], and Yiorgos Vassilandonakis[2]

[1] Computer Science Department, College of Charleston,
66 George Street, Charleston, SC 29424, USA
manarisb@cofc.edu, dsjohnso1@g.cofc.edu
[2] Music Department, College of Charleston,
66 George Street, Charleston, SC 29424, USA
vassilandonakisy@cofc.edu

Abstract. We present Harmonic Navigator (HN), a system for navigating and exploring harmonic spaces extracted from large musical corpora, to be used in music composition and performance. A harmonic space is a set of harmonies (chords) and transitions between harmonies found in a music corpus. By navigating this space, the user can derive new harmonic progressions, which have correct voice leading. HN is controllable via a Kinect gesture interface. To aid the user, the system also incorporates stochastic and evolutionary techniques. HN offers for two primary modes of interaction: a harmonic transition selector, called *harmonic palette*, which utilizes a GUI to navigate harmonic transitions in a front-to-back manner; and a harmonic-flow scrubber, which presents a global overview of a harmonic flow and allows the user to perform common audio scrubbing and editing tasks. Both GUIs use colors to indicate harmonic density based on Legname's density degree theory.

Keywords: harmonic navigation, computer music, graphical user interface, gesture language, Kinect sensor, harmonic space, music composition, music performance.

1 Introduction

The use of computation in music composition and performance has emerged from advancements in music technology, such as MIDI interface and synthesized instruments, explorations in the use of mathematic and aleatoric principles in composition by composers like Iannis Xenakis, Gyorgy Ligeti and John Cage [1,2], and the application of artificial intelligence tools to music analysis and generation.

Several systems have emerged in recent decades to assist with music performance and composition, including Cope's *EMI* [3], Biles' *GenJam* [4], and Pachet's *Continuator* [5]. These are discussed in more detail in the next section. We present a novel system that provides an innovative, gesture-driven user interface for navigating

M. Kurosu (Ed.): Human-Computer Interaction, Part II, HCII 2014, LNCS 8511, pp. 58–68, 2014.

harmonic spaces of music from large corpora. This system combines stochastic and evolutionary techniques and is an extension of *Monterey Mirror*, an interactive system for melodic exploration [6].

Harmonic Navigator allows for user-guided generation of new harmonic (chord) material from an existing musical corpus (currently we explore the Riemenschneider collection of *371 J.S. Bach Chorales*). This corpus is used to train a Markov model, a stochastic model that represents the transition probabilities of chords in the corpus. The Markov model is capable of rapidly generating material that is similar to the provided corpus. In practice, the generated material often contains only short-term similarities (event-to-event) and lacks long-term coherent structure. We utilize a genetic algorithm to search the Markov model for high-quality material. Using power-law metrics as a fitness measurement allows the genetic algorithm to select material that is similar to target material, such as a user-provided melody or harmonic flow [7-8].

Finally, the system allows saving of a generated chord progression, for further processing and use in music composition projects.

This paper focuses on the user interface aspects of the Harmonic Navigator. The remaining sections are organized as follows: section 2 presents related background research; section 3 describes the target audience and presents a high-level task analysis for the system; section 4 describes the user interface in more detail; section 5 provides an overview of the system architecture and major software components; finally, section 6 discusses future work.

2 Background

Within the last 50 years there have been numerous applications of computing and artificial intelligence in analysis, generation, composition, and performance of music. While these results are sometimes judged by how well they model human intelligence (strong AI), the real contribution lies in how they may enhance human creativity and facilitate artistic exploration and expression.

GenJam generates jazz improvisations for real time performance [4]. *GenJam* is trained using an interactive genetic algorithm, which determines fitness through a human mentor. The trained population is used to "trade fours" with a human performer.

The *Corpus-Based Harmonic Progressions Generator* [10] mixes stochastic selection via Markov models and user input to generate harmonic progressions in real time. The user enters information to specify harmonic complexity and tension, as well as a bass-line contour. This is used by the system to influence the selection of harmonies from the trained Markov models, and to generate a harmonic progression.

Experiments in Music Intelligence (EMI) is the most comprehensive work in automated computer music generation to-date [3]. EMI analyzes a corpus of musical works and trains Markov models. EMI can then automatically compose pieces in a style similar to the corpus. EMI works offline and has been used to generate numerous pieces in the style of various composers.

Continuator is an interactive music performance system which accepts musical input from a human performer. It completes musical material in the same style as the user input [5]. Using a musical corpus, the system trains several Markov models. Human performer input is matched against the various Markov models until a match is found. The corresponding Markov model is used to generate a musical continuation. This makes the system sometimes generate perfect reproductions of earlier musical input, and other times less accurate repetitions (introducing interesting variations).

NEvMuse [11] is an experiment in using genetic programming to evolve music pieces based on examples of desirable pieces. *NEvMuse* uses power-law metrics as fitness functions. In an evaluation experiment, these metrics were able to predict the popularity of 2000 musical pieces with 90.7% accuracy. The system was used to autonomously "compose" variations of J.S. Bach's Invention #13 in A minor (BWV 784). Similarly to *NevMuse*, the Navigator's genetic algorithm uses power-law metrics to determine fitness.

Monterey Mirror [6] uses Markov models, genetic algorithms and power-law metrics to generate musical phrases in real-time, based on musical input from a human performer. Markov models are used to capture short-term correlations in melodic material. The genetic algorithm is then used to explore the space of probable note combinations, as captured by the Markov model, in search of novel, yet similar melodic material. Similarity is measured using power-law metrics, which capture long-term correlations in melodic material, i.e., the statistical balance between expectation and surprise across various musical parameters [8].

Harmonic Navigator is implemented in Jython and Java using custom GUI, MIDI and OSC libraries. It incorporates computational elements from *NevMuse* and *Monterey Mirror* to allow human performers to navigate the space of musical harmonies using a gesture-based interface [12].

In this paper, we present a new user interface for the Harmonic Navigator that allows composers and performers to create new music by modifying musical output of a generative system in real-time.

3 Target Audience

The Harmonic Navigator (HN) is a gesture-based interactive system for exploring harmonic spaces of distinct (or composite) musical styles (see Fig. 1). Also, it may be used to generate music in real-time, in collaboration with human performers. The UI has been designed for users with, at least, basic training in functional tonality (first-year college harmony, or equivalent). However, as we collect usability feedback from more users, this UI may evolve (e.g., to provide more or, even, less information).

Fig. 1. One of the authors interacting with the system

3.1 Music Composers

HN can be incorporated by music composers, in a computer-aided composition context. In particular, a composer may use it to explore compositional ideas in harmonic spaces derived from various musical corpora. These can consist of pre-existing libraries of established musical (and therefore harmonic) styles, or could be a collection of the composer's previous own body of work. By employing these corpora, traditional harmonies may be derived and be evaluated on a consonance/dissonance scale [14]. More dense harmonies may also be explored, and may be similarly evaluated on a consonance/dissonance scale [9]. In order for this to work well, the music corpora loaded to the system must contain enough musical pieces (for harmonic variety) and should be stylistically consistent (e.g., consist only of Baroque pieces, or Impressionist pieces). By combining two stylistically inconsistent groups of pieces, this would create a rather disjoint harmonic space, consisting of two mostly isolated "islands" (although it would be possible to "travel" from one to the other, via, some common basic harmonies, which may appear in pieces from both styles, but function in different ways in each).

3.2 Music Educators

HN may also be used to enhance traditional classroom pedagogy in tonal harmony. Professors may engage students through tonal harmony games implemented on an HN platform. "Players" could interactively assign appropriate tonal function and hierarchy to each important pitch in a melody: tonic, predominant or dominant [13], and then select from a variety of available chords in various inversions. In the end, users may gain a much deeper appreciation for harmonic functions quicker and retain it for a much longer time.

It could also be used to explore pitch-set generated harmony or 12-tone and serial harmonic styles in a more advanced 20th century harmony course. In this context, HN will provide even more insight to the student, as it would be able to offer suggestions that take into account pitch usage and cycling, in addition to the harmonic spacing and density.

3.3 Music Performers

Finally, HN may be used in musical performances. Musicians and non-musicians (e.g., members of the audience, or passers-by), may utilize MIDI and OSC controllers (e.g., iPhone TouchOSC client), as well as traditional instruments, to create harmonic contexts for improvised performances. Another related possibility is to compose/design musical games (e.g., the system could be driven through audience participation) to engage, inspire, and possibly challenge musicians in various performance environments, or to allow non-musicians to create musical performances in unconventional settings (such as art galleries with HN sound installations).

4 User Interface

The Harmonic Navigator offers two primary modes of interaction: a gesture-based harmonic transition selector, called the *harmonic palette*, and a harmonic-flow scrubber, which presents a global view of a flow being generated. The first UI provides a tree-level view, and thus allows for localized control and inter-harmony navigation. The second UI provides a forest-level view, and supports scrubbing and editing actions. Herein we focus mainly on scrubbing actions (such as playback in arbitrary speed). Both views use colors to indicate harmonic density calculated using Legname's density degree theory [9].

4.1 The Harmonic-Palette View

The Harmonic-Palette View presents available harmonies as a dynamic navigable space. It utilizes a 3D front-to-back approach. The interface presents users with a *harmonic palette*, from which to choose a follow-up harmony (see Fig. 2). The palette contains a number of circles, each representing a harmony. The current harmony is located in the center of the display. Follow-up harmonies are determined by the current harmony (as dictated by the training corpus), and are placed in a clockwise fashion, around a clock face labeled with the 12 tones. Pieces are normalized to the tonic, so pitch C is always positioned at 0. We use vertically stacked numbers to denote harmonic intervals. This is consistent with the vertical placement of notes on a staff. We have considered using Western musical notation, however, this representation provides more direct information, i.e., users can see the intervals right away - they do not have to derive them from the musical notation.

Moreover, the size (radii) of follow-up circle-harmonies corresponds to transition probabilities from the current harmony (the larger, the more probable).

In the case of multiple follow-up harmonies having the same root pitch (e.g., see E and A root pitches, in Fig. 2, they are arranged around a smaller clock face. The size (radius) of this clock face corresponds to the sum of the enclosed harmonies' probabilities. Hovering the cursor over this clock face zooms in to display a larger version of the clock face, which presents more information about the contained harmonies, and allows the user to select one. These harmonies are arranged inside the smaller clock face based on the second pitch in the harmony.

When dealing with multiple harmonies that have the same second pitch, these harmonies will also be placed inside an even smaller clock. This hierarchical grouping continues until all harmonies can be represented individually.

The HN engine is capable of making recommendations for what it considers possibly aesthetic choices for follow-up harmonies. This is accomplished via a genetic algorithm which runs continuously (in the background) to suggest interesting harmonic flows. The follow-up harmony (or harmonies) selected by the genetic algorithm is (are) identified by a special bright ring around a suggested harmony. Since the genetic algorithm is running continuously, it is possible for the suggested harmony to change (by the genetic algorithm discovering a better choice) as the user is contemplating.

Circle-harmonies are assigned color based on intervallic tension. Since intervallic tension is already visible on the interface, through the displayed harmonic intervals, the assigned color representation is redundant. This emphasizes the existing information, and makes it more visible to non-experts.

Intervallic tension of a chord is determined by two factors. One is the intervallic content of the chord - a chord with more tense intervals has a higher tension factor, and thus sounds more dissonant. The relaxation vs. tension of the chord is mapped to cool vs. warm colors on a color wheel, i.e., blues are cool (relaxed) and reds or yellows are warm (tense).

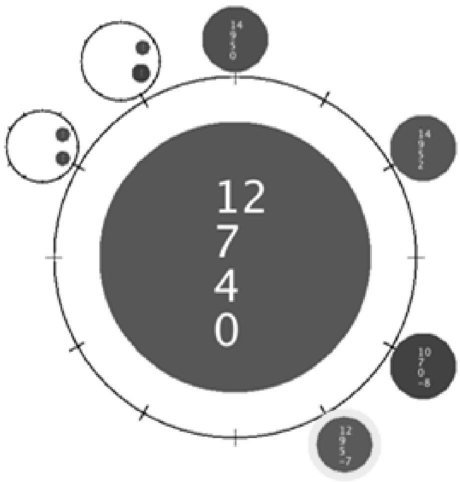

Fig. 2. The *harmonic palette* interface is used to select a follow-up harmony. The current harmony is in the center. *Follow-up harmonies* are arranged in a clockwise fashion, around a *clock face* corresponding to the 12 tones. *Numbers* represent harmonic intervals in a chord. *Color* (here reproduced in grayscale) denotes chord harmonic density.

4.2 Gesture Language for the Harmonic Palette

The Harmonic Palette UI has been designed to support three main user tasks for harmonic navigation. These are:

1. "Explore the current harmony palette";
2. "Select a follow-up harmony"; and
3. "Backtrack" (i.e., unselect current harmony and return to the previous palette).

Our current prototype is implemented using the Kuatro system. The Kuatro system is a new architecture for supporting a multitude of sensors and wireless controllers for audio/visual interactive installations. The main objective behind its design is to hide the complexities of communicating with such devices, and allow the UI developer to focus on the higher-level aspects of designing an effective UI for audio/visual control of a computer music application. The Kuatro architecture will be reported elsewhere.

We have designed a Kinect-based gesture language to implement the above user tasks. (We are also exploring gesture languages for other controllers, such as the Leap Motion sensor and OSC control via smartphones.) The Kinect gesture language utilizes only one hand via three gestures (freeing the second hand for other activities, such as interacting with MIDI and OSC controllers):

- **Hand Movement in the X-Y Plane** – Moving the hand left-to-right and up-to-down moves the cursor around the display. This action supports exploration of the current harmony palette (e.g., hovering over a secondary clock face to enlarge it).
- **Hand Push** – Pushing towards a follow-up harmony selects it. This moves the selected circle-harmony to the center, begins sounding the corresponding harmony (via MIDI), and displays the next harmony palette. This action supports moving forward in the harmonic space.
- **Hand Wave** – Waving over the current circle-harmony (center of the display), stops sounding it, and returns to the previous harmonic palette (to possibly try something else). This action supports moving backward in the harmonic space.

In particular, moving backwards allows the user to step back to previous harmony selection points, and try other alternatives. While this may seem peculiar during live performance, it may be utilized creatively (not unlike sound looping, and/or "scratching" by DJs). On the other hand, this is quite natural for composition tasks (i.e., "should I use this harmony or that?" or "what harmonic choices would I have here, had I gone to a relative minor three chords ago?").

4.3 Harmonic-Flow View

The Harmonic-Flow View presents a global, forest-level view of a harmonic flow generated by the user through the harmonic palette UI (or automatically by the harmonic generator engine). Through this view, the user may explore and update the different harmonies comprising the harmonic flow as they desire. As seen in Fig. 3, harmonies are placed horizontally across the display. For each harmony being selected, alternate harmonies, as determined by the Markov model, are displayed vertically.

Fig. 3. The *harmonic-flow scrubber* interface is used to view a complete harmonic flow, as constructed through the lower-level interface (see Fig. 2). The harmonic flow appears on the horizontal. Individual harmonies are displayed as *rectangles*. Hovering over a rectangle presents alternative harmonies (on the vertical). *Color*, again, denotes chord harmonic density.

4.4 Gesture Language for the Harmonic Flow

The Harmonic Flow UI supports three main user tasks. These are:

1. "Forward and backward scrubbing";
2. "Explore alternative harmonies"; and
3. "Select a new harmony".

As mentioned earlier, through the Kuatro architecture, users may utilize various gesture and motion controllers to interact with the UI. Herein, we present a Kinect-based language for users to control the system via their location in a room (many other possibilities exist for other controllers and sensors). By viewing the room from above, we use an X-Y coordinate system to track a user through the room and map their location to specific tasks.

- **Movement along the x-axis** – The x-axis runs parallel to the display and controls the scrubbing capabilities. By moving parallel to the display, the user identifies which harmonies are played across the flow. The tempo of scrubbing is controlled by how fast or slow the user is moving in this direction.
- **Movement along the y-axis** – The y-axis is perpendicular to the display. By moving along the y-axis, the user plays the harmonies presented in the vertical list of harmonic alternatives. As the user moves closer to the display, they play harmonies upward in the transition list; these are harmonies with increasing tension. As the user moves away from the display, they play harmonies downward in the transition list; these are harmonies with decreasing tension.

A user selects an alternative harmony by beginning to move again across the x-axis. Also, selecting an alternative harmony triggers HN to regenerate the flow based on their new selection, if opted by the user, via the genetic algorithm. The genetic algorithm and the corresponding automated generation of harmonic flows is presented extensively in [15].

5 System Architecture

The Harmonic Navigator system uses a Model-View-Controller architecture based on the Kuatro system. This reduces complexity from the UI design and implementation, while allowing for a multitude of controllers, such as a mouse, a Kinect, and smartphones using OSC clients (e.g., TouchOSC). In Fig. 4, the View in this architecture is the UI, the controller is the Gesture Engine, and the Model is the Harmonic Generator. To support a wide range of controllers we have implemented a protocol for the Gesture Engine to communicate via OSC. (This will be presented in a future publication.)

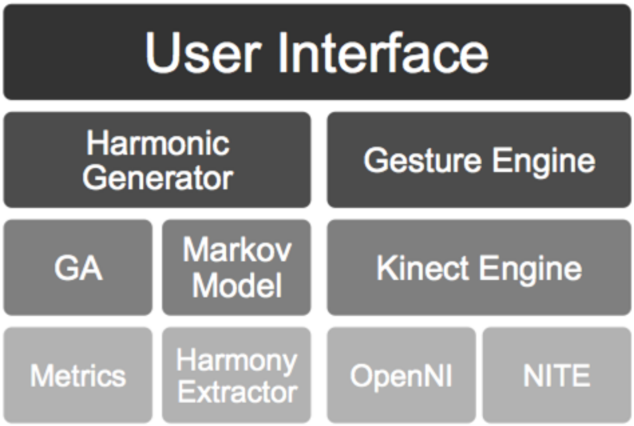

Fig. 4. Harmonic Navigator architecture using a Microsoft Kinect

6 Discussion and Future Work

The Harmonic Navigator is a powerful tool for exploring harmonic spaces in a direct, physical, and accessible manner. As new gesture control devices are introduced, its power will only increase. The possibility of allowing non-expert musically users to experience harmonic flows in such an intimate manner presents various possibilities for further work. We are currently exploring an application that will introduce the novice theory student to the notion of tonal function in common practice music. The system attaches a T, PD or D label to each suggested chord in the *harmonic flow scrubber*, and the user can quickly develop their listening ability to recognize tonal function and navigate harmonically through a musical phrase using harmonic implications alone. They can then harmonize a given melody or bass line in a more musically intelligent way by selecting chords with the appropriate function label among the ones suggested by HN, thus gaining a deeper understanding of tonal harmony. This deeper understanding would normally take several years of study, as well as keyboard proficiency. This type of learning could be further enhanced by creating a physical space larger than the user, so that he can navigate through it by walking around the space, "scrubbing" through the functional harmonic space, as possible via the Kuatro architecture discussed above.

We have presented Harmonic Navigator, a system for navigating and exploring harmonic spaces extracted from large musical corpora, to be used in music composition and performance. This system is currently being evaluated with actual users, in order to improve its usability and possibly improve its UI.

In closing, video demos of the system are available here:

- A demo of the harmonic palette UI being controlled via a Kinect - **http://goo.gl/ni7ZVl**.
- A demo of the harmonic flow view - **http://goo.gl/hpXk2G**.

References

1. Xenakis, I.: Formalized Music: Thought and Mathematics in Music. Pendragon Press, Hillsdale (1992)
2. Cage, J.: Silence: Lectures and Writings of John Cage. Wesleyan University Press, Middletown (1961)
3. Cope, D.: Virtual Music: Computer Synthesis of Musical Style. MIT Press, Cambridge (2004)
4. Biles, J.A.: Performing with Technology: Lessons Learned from the GenJam Project. In: Musical Metacreation Worksop, 9th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2013), pp. 14–19. AAAI Press, Palo Alto (2013)
5. Pachet, F.: Playing with Virtual Musicians: The Continuator in Practice. IEEE Multimedia 9(3), 77–82 (2002)
6. Manaris, B., Hughes, D., Vassilandonakis, Y.: Monterey Mirror: Combining Markov Models, Genetic Algorithms, and Power Laws. In: 2011 IEEE Congress on Evolutionary Computation (CEC 2011), 1st Workshop in Evolutionary Music, pp. 33–40. IEEE Press, New York (2011)
7. Zipf, G.K.: Human Behavior and the Principle of Least Effort. Hafner Publishing Company, New York (1949)
8. Manaris, B., Romero, J., Machado, P., Krehbiel, D., Hirzel, T., Pharr, W., Davis, R.B.: Zipf's Law, Music Classification and Aesthetics. Computer Music Journal 29(1), 55–69 (2005)
9. Legname, O.: Density Degree of Intervals and Chords. 20th Century Music 4(11), 8–14 (1997)
10. Eigenfeldt, A., Pasquier, P.: Realtime Generation of Harmonic Progressions Using Controlled Markov Selection. In: 1st International Conference on Computational Creativity (ICCC-X), pp. 16–25. ACM Press, New York (2010)
11. Manaris, B., Roos, P., Machado, P., Krehbiel, D., Pellicoro, L., Romero, J.: A Corpus-Based Hybrid Approach to Music Analysis and Composition. In: 22nd Conference on Artificial Intelligence (AAAI 2007), pp. 839–845. AAAI Press, Palo Alto (2007)

12. Manaris, B., Johnson, D., Vassilandonakis, Y.: Harmonic Navigator: A Gesture-Driven, Corpus-Based Approach to Music Analysis, Composition, and Performance. In: 9th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2013), pp. 67–74. AAAI Press, Palo Alto (2013)
13. Berry, W.T.: Structural Functions in Music, pp. 40–57. Prentice Hall, Upper Saddle River (1976)
14. Hindemith, P.: The Craft of Musical Composition, Schott, Mainz, Germany, pp. 87–89 (1945)
15. Manaris, B., Johnson, D., Vassilandonakis, Y.: A Novelty Search and Power-Law-Based Genetic Algorithm for Exploring Harmonic Spaces in J.S. Bach Chorales. In: 3rd International Conference on Evolutionary and Biologically Inspired Music, Sound, Art and Design. Springer, Berlin (2014)