An evolutionary approach to algorithmic composition

JONATAS MANZOLLI,† ARTEMIS MORONI,‡ FERNANDO VON ZUBEN§ and RICARDO GUDWIN||

† Interdisciplinary Nucleus of Sound Communication (UNICAMP/NICS) and §|| Electrical Engineering Faculty (UNICAMP/FEE), University of Campinas, Rua da Reitoria, 165-13081-970, Campinas SP, Brazil
‡ The Automation Institute (CTI/IA), Technological Center for Informatics, Via Dom Pedro I, KM143/6, Campinas SP, Brazil
E-mail: † jonatas@nics.unicamp.br, ‡ artemis@ia.cti.br,
§ vonzuben@dca.fee.unicamp.br, || gudwin@dca.fee.unicamp.br

This paper presents an end-user interface that allows realtime parametric control of sound events resulting in an interactive environment, in which evolutionary computation is applied to algorithmic composition. The resulting system, Vox Populi, uses genetic algorithms to generate and evaluate a sequence of chords played as MIDI data. Harmonic, tonal and voice range fitness are used to control musical features. Based on the ordering of consonance of musical intervals, the notion of approximating a sequence of notes to its harmonically compatible note or tonal centre is used. This method employs fuzzy formalism and is posited as an optimisation approach based on factors relevant to hearing music.

1. INTRODUCTION

Evolutionary computation has been successfully applied to control music processes. An application of genetic algorithms to generate jazz solos is described in Biles (1994) and this technique has also been studied as a way of controlling rhythmic structures (Horovitz 1994).

This paper describes a procedure for algorithmic composition based on the controlled production of chord cadences: a population of chords is properly codified according to the MIDI protocol, and is submitted to evolution by the application of genetic algorithms. The research presented here has already been discussed in Moroni, Manzolli, Von Zuben and Gudwin (1999). A fitness criterion is defined to indicate the best chord at each generation, and this chord is selected as the next element in the sequence to be played. The produced realtime sound output stream allows this system to be used in live electronic music as a computer performance instrument.

In what follows, a general description of the main components of the computational environment are presented and a musical fitness based on melodic, harmonic and voice range criteria is defined. Finally, the end-user interface is described in detail.

2. DEFINITIONS

The following basic concepts are fundamental to the understanding of the main results.

2.1. Individuals and populations

Individuals of the population are defined here as chords of four notes and they are potential solutions for the selection process. Initially, the chord's notes are randomly generated on the interval [0 . . . 127], which corresponds to the range of MIDI note numbers. In each generation, a population of thirty chords is produced and evaluated. The chords are internally represented as a chromosome with 28 bits, composed of 4 words with 7 bits each (figure 1).

2.2. Rhythmic genetic cycle

The general architecture of the rhythmic genetic cycle is depicted in figure 2. It is possible to see on the diagram two cooperative processes in the genetic cycle: one producing notes and the other (the interface) consuming notes. As described in section 2.1, once the initial population of individuals is created, the fitness of each chord can be evaluated. The fitness function is defined as a composition of three terms: voice range fitness, vertical consonance or harmonic fitness, and horizontal consonance or tonal fitness.

After the music fitness evaluation, typical genetic operators of crossover and mutation are applied to the individuals (Michalewicz 1996). Once the best chord is selected, it is available to be played. The interface, which is looking for new notes, sends them to the MIDI port.

The following steps are realised in the genetic cycle (Pedrycz and Gomide 1998):

(1) Create an initial population randomly.
(2) While not stopped (by the user), perform the following:

- Evaluate the musical fitness of each individual in the population.

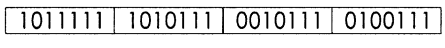


Figure 1. The structure of a MIDI chromosome.

https://doi.org/10.1017/S1355771899002083 Published online by Cambridge University Press

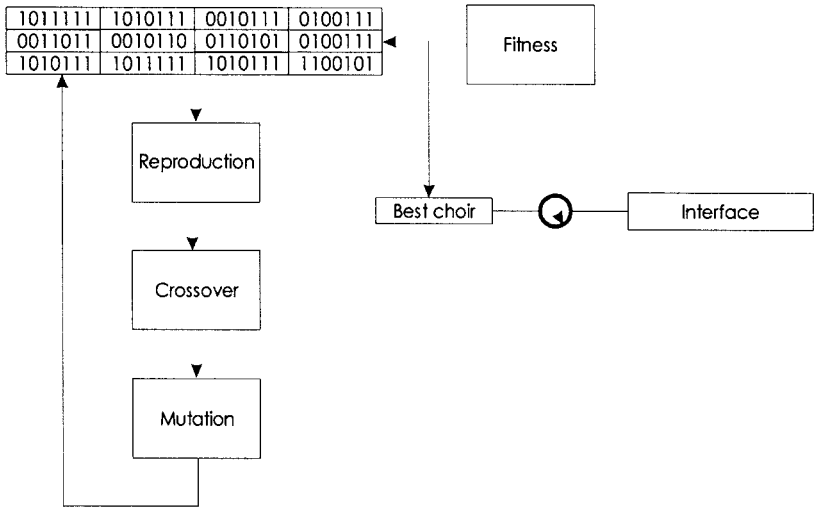


Figure 2. The rhythmic genetic cycle.

- Apply the genetic operators to the population of MIDI chromosomes (groups of voices), taken into account based on the musical fitness, to create a new population. That is:
  - (i) *Reproduction*. Copy existing individual strings to a new population.
  - (ii) *Crossover*. Create two new chromosomes by crossing over randomly chosen sublists (substrings) from two existing chromosomes.
  - (iii) *Mutation*. Create new chromosomes from an existing one by randomly mutating the character in the list.

(3) Find the best individual in the new population and play it as a MIDI event. Go to step 2.

The steps above stress the existence of many operations executed in each cycle. The time interval between the selection of the best chords in two successive cycles may be different. On the other hand, the interface is regularly 'asking for new notes'. Despite the fact that there is an average time cycle to designate the best individual in each generation, small variations in each time cycle determine the *genetic rhythm*. Different times for the notes being played are perceived as a rhythmic profile of the pitch sequence generated by the genetic cycle.

3. FITNESS EVALUATION

This section presents the criteria used to evaluate the system musical fitness associated with each chord. Mathematical formulation is omitted; a detailed description is found in Moroni *et al.* (1999).

3.1. Voice range criterion

The chord's notes are related to voices that are associated with the linguistic terms: bass, tenor, contralto, soprano and nh (no human). The related fuzzy sets are shown in figure 3. Each voice is assigned to a membership value associated with each linguistic term in the set {NH, B, T, C, S}. For the classification of each voice, the membership function is evaluated for each set and the maximum value is taken. In the case of coincidence, the distance to the centre of the fuzzy set is considered. The interval of voices reached by the human voices is assumed to be in the interval $H = [39 \ldots 84]$, given in MIDI note values.

Once the voices of each chord are evaluated according to their distribution in terms of intervals, the voice range criterion returns a value in the set {NH, W, M, G, E}. These linguistic values are associated with the concepts: *No Human*, *Weak*, *Medium*, *Good* and *Excellent*. The optimal case – Excellent – is considered when the chord contains the voices: bass, tenor, contralto and soprano; in this case, Nvalues = 4. The absence of these voices returns NH, the presence of one of them returns W, two returns M, three returns G, with Nvalues = 0, 1, 2, 3, respectively. Therefore, the voice range fitness is evaluated as Nvalues/4.

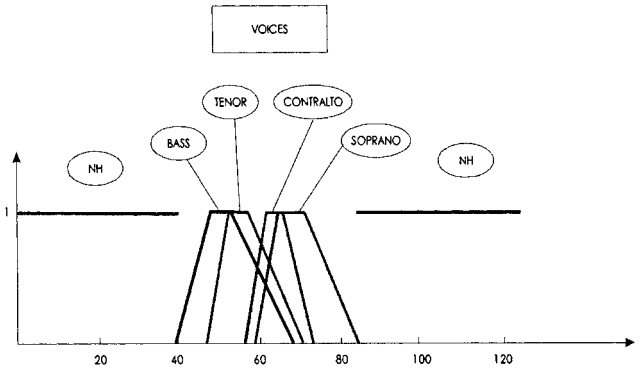


Figure 3. The linguistic values associated with the voices.

https://doi.org/10.1017/S1355771899002083 Published online by Cambridge University Press

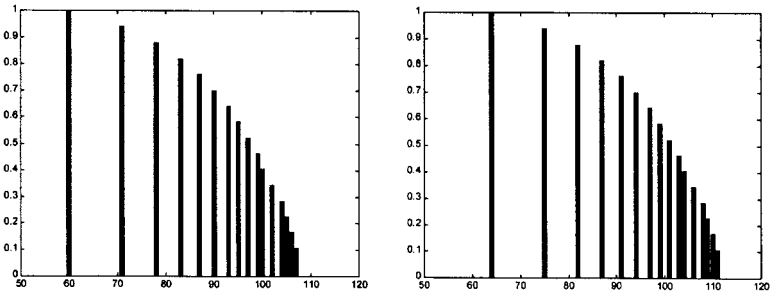


Figure 4. The weighting of notes 60 (C) and 64 (E).

3.2. The consonance criterion

The consonance among the four voices is evaluated as a function of the voice attributes. Consonance is defined as a function of the overlap between the harmonic series of two given notes (Vidyamurthy and Chakrapani 1992). This overlap measurement is then scaled to a value between 0 and 1, with 1 denoting complete overlap (i.e. the two notes being the same), and 0 denoting no overlap at all. This notion of overlap can be succinctly captured in the fuzzy set formalism described here.

The harmonic series derived from a given note is a set tone consisting of its fundamental tone and upper harmonic tones. In figure 4, the weighting of the harmonic series versus the relative pitch is represented, and the sum of the weights is normalised. Note that n denotes the nth key on the piano, and that $(n+k)$ denotes the key k semitones above key n. In figure 4, the upper tones of notes 60 (C) and 64 (E) are presented.

Formally defined, each note is a fuzzy set on a countable universe of discourse. The consonance or overlap between the notes is defined as the sum of the intersection of the harmonic series weights (Vidyamurthy and Chakrapani 1992), and results in a value in $[0 \ldots 1]$. Figure 5 shows the consonance between the pitches shown in Figure 4.

3.2.1. Vertical consonance or harmonic fitness

Given a four-note chord, the harmonic fitness is defined as the sum of the consonance or overlap of the harmonic series of the four notes present in the chord (this is shown for two notes in figure 5).

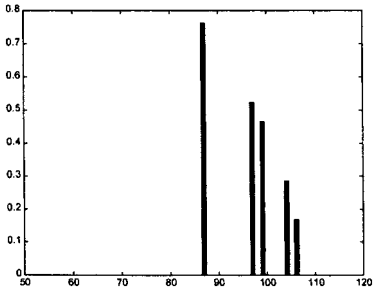


Figure 5. The consonance or overlap between note 60 (C) and note 64 (E).

3.2.2. Horizontal consonance or tonal fitness

Given a four-note chord and a tonal centre named Id, the tonal fitness is defined as the maximum value of the consonance between each note present in the chord and the tonal centre.

3.3. Musical fitness

The resulting musical fitness is a conjunction of the previous functions and is defined as

$$\text{Musical fitness} = \text{voice range fitness} + \text{harmonic fitness} + \text{tonal fitness}.$$

4. VOX POPULI SYSTEM

The system was implemented to perform a series of sound experiments and eventually to be used as a tool for algorithmic composition.

4.1. Interface and parametric control

The user may interfere in the fitness function through five interface controls: (i) tonal centre control, (ii) biological control, (iii) rhythmic control, (iv) voice range control, and (v) orchestra control. There follows a short description of the controls available for user interaction with Vox Populi (see figure 6).

4.1.1. Tonal centre control

The Mel scroll allows modification of the value Id, which is the tonal centre in the evaluation of the tonal fitness.

4.1.2. Biological control

The Bio scroll allows interfering in the duration of the genetic cycle, modifying the time between genetic iterations. Since the music is being generated in real time, this artifice was necessary to synchronise the

https://doi.org/10.1017/S1355771899002083 Published online by Cambridge University Press

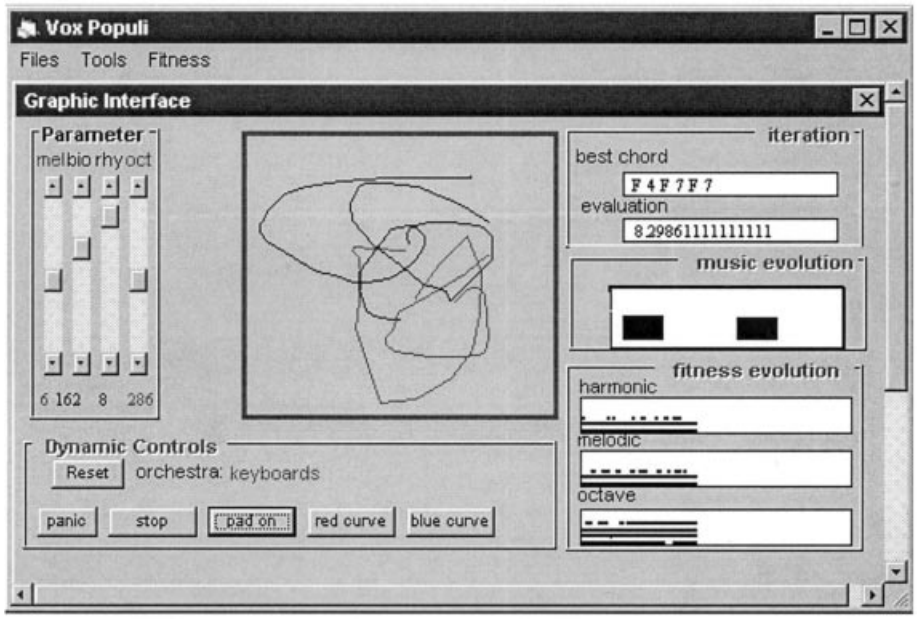


Figure 6. Vox Populi graphic interface.

different processes that are running. This value determines the slice of time necessary to apply the genetic operators, such as crossover and mutation, and can also be interpreted as the reproduction time for each generation.

4.1.3. Rhythmic control

The Rit scroll changes the time between evaluations of the musical fitness. It determines the 'time to produce a new music generation' or the slice of time necessary to evaluate the musical fitness of the population. It interferes directly with the musical rhythm. Changes of this control make the rhythm faster or slower.

4.1.4. The voice range control

The Oct scroll allows enlarging or diminishing the interval of voices considered in the voice range criterion.

4.1.5. The orchestra control

Six MIDI orchestras are used to play the selected chords: (i) keyboards, (ii) strings and brass, (iii) keyboards, strings and percussion, (iv) percussion, (v) sound effects, and (vi) random orchestra, which takes any instrument from the General MIDI list.

4.2. Interactive pad control

The Pad On button enables and disables the pad change on four controls defined above. They are coupled in two pairs that are interpreted as variables of a two-dimensional phase space. This allows a user to draw an oriented curve to determine the musical evolution. There are two curves, associated with different colours:

(1) *Red curve*. This describes a phase space of the *tonal* and *voice range* control variables.
(2) *Blue curve*. This describes a phase space of the *biological* and *rhythmic* control variables.

The pad may be musically interpreted as an elementary tool that allows a 'master gesture' to conduct the music.

5. PERFORMANCE CONTROL

The Performance Control window shown in figure 7 allows the user to shape the musical output in real time. Each voice can be played as: (i) Solo, (ii) Arpeggio, and (iii) Chord. The Solo option enables direct output to the MIDI port. The Arpeggio option uses the last four notes calculated by the algorithm as MIDI output. The Chord option plays the last four

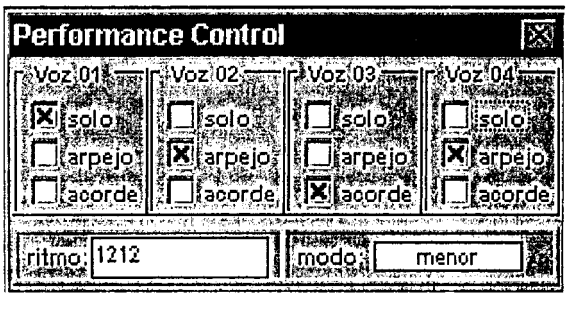


Figure 7. Performance Control options.

https://doi.org/10.1017/S1355771899002083 Published online by Cambridge University Press

notes as a block. A musician can modify pitch material to chromatic, major or minor modes using the Modo Option. Rhythmic Control, presented in section 4.1.3, can be shaped by a rhythmic pattern input by the user as a string of small integer numbers. Vox Populi has also a Write MIDI option to produce MIDI files to be processed later in a sequencer program. A Vox Populi demo can be downloaded from http://www.ia.cti.br/~artemis/voxpopuli

6. CONCLUSION

The music resulting from the Vox Populi system moves from very pontilistic sounds to sustained chords, depending on the duration of the genetic cycle and the number of individuals in the original population. The octave fitness forces the notes to be in the range H, assumed to be the range covered by the human voice and associated with the central region of the notes on the piano. However, since several orchestras of instruments are used, this range is too limited for some of them. The original decision to restrict the generated voices to specific ranges was just to resemble human voices; nevertheless, a user can enlarge this range using the Octave Control. The interface was designed to be flexible for the user to modify the music being generated, and can be thought of as a prototype environment for algorithmic composition. The system has proven to be adaptable and to be able to receive new features. Most of the controls did not exist in the first version, and were added as new interesting aspects that were identified for extension. It has shown itself to be a laboratory of sound experiments in which many possibilities can be explored.

Further, we are developing the integration of this system with gesture interfaces, such as a glove to enhance the man–machine interaction, with the purpose of allowing a human gesture to be the realtime controller. This is a natural extension of the Pad Control. This approach was previously applied in the project ActContAct (Manzolli, Moroni and Matallo 1998), where an electronic shoe was used in music generation.

REFERENCES

Biles, J. A. 1994. GenJam: a genetic algorithm for generation of jazz solos. *Proc. of the Computer Music Conf. (ICMC '94)*, pp. 131–7.

Horovitz, D. 1994. Generating rhythms with genetic algorithms. *Proc. of the Computer Music Conf. (ICMC '94)*, pp. 142–3.

Manzolli, J., Moroni, A., and Matallo, C. 1998. AtoCon Ato: new media performance for video and interactive tap shoes music. *6th ACM Int. Multimedia Conf.*, p. 3.

Michalewicz, Z. 1996. *Genetic Algorithms + Data Structures = Evolution Programs*, pp. 283–7. Springer-Verlag.

Moroni, A., Manzolli, J., Von Zuben, F., and Gudwin, R. 1999. Vox Populi: evolutionary computation for music evolution. *Proc. of the AISB '99 Symp. On Creative Evolutionary Systems*, pp. 82–5. ISBN 1902956036, Published by AISB.

Pedrycz, W., and Gomide, F. 1998. *An Introduction to Fuzzy Sets Analysis and Design*. Cambridge, MA: MIT Press.

Vidyamurthy, G., and Chakrapani, J. 1992. Cognition of tonal centers: a fuzzy approach. *Computer Music Journal* **16**(2). Cambridge, MA: MIT Press.

https://doi.org/10.1017/S1355771899002083 Published online by Cambridge University Press