Learning large scale musical form to enable creativity

Francis Screene and Geraint A. Wiggins
Computational Creativity Lab
School of Electronic Engineering and Computer Science
Queen Mary University of London
Mile End Road, London E1 4FZ, UK
fscreene@icloud.com,geraint.wiggins@qmul.ac.uk

Introduction

We present a preliminary study of a novel method for the creation of original music from learned data. Historically, the vast majority of music creation programs have worked from the bottom up, constructing music from notes and/or chords, according to music-theoretic or machine-learned rules (e.g. Ebcioğlu, 1988; Cope, 1991; Conklin and Witten, 1995; Pearce and Wiggins, 2007). Our approach is different. We are interested in the high-level structure of music, rather than the very low level. Starting from work by Collins (2011), we attempt to merge stochastic methods for music generation with top-down analytical methods for music-structural analysis (Meredith, Lemström, and Wiggins, 2002). The long term vision behind our proposal is to learn a statistical structural grammar from collections of pieces, which can then be used to generate large-scale structures which in turn can be instantiated with low-level detail. In this paper, we take the very first step: to analyse a piece of music into a graph of structural components and use it to produce music which might be considered as an improvisation on that original piece.

Background

SIA

The SIA algorithm was developed by (Meredith, Lemström, and Wiggins, 2002). The objective of the algorithm is to exhausitvely identify repeating patterns and structure within a given point set, used, in our case, to represent a musical score.The algorithm allows automatic extraction of such repeated structural units as shown in the two boxes in Figure 1, and a concise representation of the relationships between each unit and its repetition, which may be at a different pitch. In general, the structural units discovered may span very small and very large segments of musical time. They are often skeletal forms, around which musical ornamentation is draped: in this case, SIA discovers the repeating skeleton and not the ornamentation. Importantly, the units that SIA discovers sometimes contain other, smaller units. This quasi-recursive structure presents interesting challenges in developing our technique.

SIA works because repetition (usually with variation) is what gives music its sense to a listener. It is this feature that allows us to use the resulting analysis for generation of new music. It yields an analysis of the piece in terms of fragments of score (called maximal translatable patters, MTPs), coupled with vectors representing all the occurrences of the same music, at whatever pitch, in the piece. These units, of one MTP plus all its translational vectors, are called translational equivalence classes (TECs). Our prototype uses a particular version of SIA, called COSIATEC, which compresses music by choosing the biggest TEC, removing and recursing until no notes are left. Then each TEC is analysed recursively also.

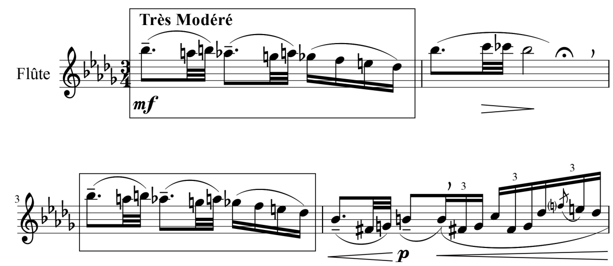

Figure 1: A very simple example of the kind of musical repetition discovered by SIA.

Probabilistic grammars

Our generation mechanism is drawn from Markovian approaches in computational linguistics (Manning and Schütze, 1999). In language, the idea is to describe likely continuations from any given initial substring of an utterance as a distribution over the possible symbols that can occur next in sequence. With some additions, this turns out to be a remarkably good model of human melodic expectation at the individual note level (Pearce and Wiggins, 2006). To our knowledge, however, there has not previously been an attempt to build a Markov model in which the states denote fragments of polyphonic music, as opposed to individual notes or chords. A particular challenge for our design was the recursive nature of the SIA computation: because TECs may contain other TECs, we must use an enriched kind of Markov model, with two kinds of arc, one to express continuation, the other to express sub-structure.

Late-breaking abstract presented at The Sixth International Conference on Computational Creativity, June 2015

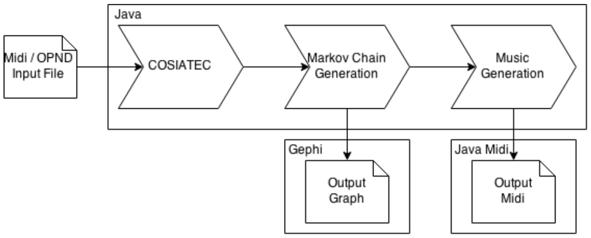

Figure 2: The architecture of our prototype.

```
Data: SIATEC Output File
Result: Markov Chain of TECs
initialization;
for Each Tec do
    CREATE events by adding vectors to TEC points;
    ADD events to eventlist;
end
SORT eventlist by timestamp;
for Each Event do
    ADD id of next events TEC to current TEC;
end
for Each TEC do
    Calculate probabilities based on TEC ids;
end
for Each TEC do
    for Each Other TEC do
        if Other TEC start >TEC start && Other TEC END >TEC END * 3 then
            Make Other TEC subtec of TEC;
        end
    end
end
```

Figure 3: The analysis algorithm.

System Architecture and Algorithms

Our prototype has a simple architecture, illustrated in Fig. 2. The learning algorithm used is shown in Fig. 3. The generation algorithms is shown in Fig. 4. The system analyses a piece of music in MIDI format, and takes a random walk around the resulting model, with probabilities determined by the structure of the SIA analysis. The result sounds somewhat like a "noodling" improvisation on the original piece.

Evaluation

We evaluated the model by generating new pieces from 10 pieces of music (2nd Time; All about that Base; Alice In Wonderland; Beale Street Blues; Maple Leaf Rag; My Favorites; Chopin Opus 22 Pt. 4; Chopin "Pathétique"; Syrinx; Toccata and Fugue in D Minor). The pieces generated from our SIA-determined model were paired in a balanced block design with uniform random walks through the same piece. The pairs were played to 16 human listeners, who were asked which of each pair sounded more intentional. The results are modest but promising: 61% chose the piece generated by our system as the more intentional one.

Conclusion and Future Work

This very preliminary study suggests that statistical learning may be a good model of musical structure at the high (section-wise) level as well as the low (note-wise) level.

```
Data: Markov Chain of TECs
Result: Sequence of Midi TECs
initialization;
currentTEC = firstTEC;
for i >maxlength do
    if currentTEC contains TECS then
        recursively call self;
    end
    else
        currentTEC = probabilistically select successor TEC from currentTEC;
        add currentTEC to playsequence;
        return playsequence;
    end
end
```

Figure 4: The generation algorithm.

However, in the best sense, it raises more questions than it answers—specifically: what is the best way to construct this non-standard hierarchical Markov Model; what is the best way to generate the new outputs; and, beyond the scope of the current work, how can multiple SIA analyses, from multiple pieces, be combined to generate music that combines ideas from multiple learned sources?

Acknowledgments

The second author was funded by the project Lrn2Cre8, which acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 610859.

References

Collins, T. 2011. *Improved methods for pattern discovery in music, with applications in automated stylistic composition.* Ph.D. Dissertation, The Open University.

Conklin, D., and Witten, I. H. 1995. Multiple viewpoint systems for music prediction. *Journal of New Music Research* 24:51–73.

Cope, D. 1991. *Computers and Musical Style*. Oxford University Press.

Ebcioğlu, K. 1988. An expert system for harmonizing four-part chorales. *Computer Music Journal* 12(3):43–51.

Manning, C. D., and Schütze, H. 1999. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: MIT Press.

Meredith, D.; Lemström, K.; and Wiggins, G. 2002. Algorithms for discovering repeated patterns in multidimensional representations of polyphonic music. *Journal of New Music Research* 31(4):321–345.

Pearce, M. T., and Wiggins, G. A. 2006. Expectation in melody: The influence of context and learning. *Music Perception* 23(5):377–405.

Pearce, M. T., and Wiggins, G. A. 2007. Evaluating cognitive models of musical composition. In Cardoso, A., and Wiggins, G. A., eds., *Proceedings of the 4th International Joint Workshop on Computational Creativity*, 73–80. London: Goldsmiths, University of London.