Evolutionary Music Composition based on Zipf's Law

Johannes H. Jensen
CRAB Lab, Department of Computer and Information Science
Norwegian University of Science and Technology
Trondheim, Norway
johannj@stud.ntnu.no

Pauline C. Haddow
CRAB Lab, Department of Computer and Information Science
Norwegian University of Science and Technology
Trondheim, Norway
pauline@idi.ntnu.no

ABSTRACT

Artificial Evolutionary techniques have shown great potential for musical tasks. One such task is the automatic generation – or composition – of music. However, when constructing evolutionary music composition systems, a major challenge is finding a suitable fitness function.

The most common fitness approach is interactive evaluation. However, efficiency challenges with such an approach has inspired the search for automatic alternatives.

In this work, a music composition system is presented for the evolution of novel melodies. The system uses an automatic fitness function based on *Zipf's Law*, which captures the scaling properties of music. Results show that the generated melodies exhibit several favourable musical properties, including melodic themes.

Categories and Subject Descriptors

F.1.1 [**Computation by Abstract Devices**]: Models of Computation—*Self-modifying machines*; J.5 [**Computer Applications**]: Arts and Humanities

General Terms

Algorithms, Experimentation

Keywords

Genetic Algorithms, Music Generation, Zipf metrics

1. INTRODUCTION

Artificial Evolutionary techniques have shown great potential for musical tasks, such as music composition. However, a major challenge for *automatic* evolutionary music composition is the design of a suitable fitness function.

An automatic fitness function is intended to steer evolution towards good music, without extensive human input. Although music exhibits many rational properties, it distinguishes itself by the focus on human emotions and aesthetics – qualities that are are not fully understood and which are difficult to describe mathematically.

Some of the approaches to musical fitness are interactive evaluation [2], hardwired rules [4] and machine learning methods [1].

GECCO'11, July 12–16, 2011, Dublin, Ireland.
ACM 978-1-4503-0690-4/11/07.

2. METHODOLOGY

Building on research on *Zipf's Law* of music [3], this paper explores the use of Zipfian metrics in automatic fitness functions for the composition of novel music.

Zipf's Law describes the *scaling properties* of many natural and social phenomena [5]. The law states that the "frequency of an event is inversely proportional to its statistical rank": $f = r^{-a}$, where f is the frequency of occurrence of some event, r is its statistical rank and a is close to 1.

Plotting the ranks and frequencies on a logarithmic scale produces a *rank-frequency distribution*, which is characterized by a straight line with a *slope* close to -1, corresponding to the exponent $-a$ in the equation. Slopes may range from zero to negative infinity, indicating uniform random to monotone distributions respectively.

Zipf metrics for music count the frequency of some *musical event* and plots them against their statistical rank on a log-log scale. Linear regression is then performed on the points to estimate the slope of the distribution.

The relevant metrics explored are: rank-frequency distributions of pitches, chromatic tones (pitch modulo 12), note durations, pitch durations, chromatic tone durations, pitch distances (time between repetitions of notes), chromatic tone distances, melodic intervals, melodic bigrams and melodic trigrams.

In this work, a fitness function based on Zipf metrics was designed which takes a pre-defined vector of *target slopes* as input. For each music individual in the population, it then calculates a fitness score based on the *distance* to the target vector. The fitness function is applied in a Genetic Algorithm for evolution of short melodies.

3. EXPERIMENTS AND RESULTS

The experiments employ a vector-based genotype, similar to that of [2], with genes representing discrete musical events (not to be confused with events in the Zipf metrics). Each event is a (type, pitch) tuple, where type is either "note" or "hold". A note event indicates the beginning of a new note, while a hold event extends the duration of a previous note. The genotype was limited to a length of 4 bars (measures) with a resolution of $\frac{1}{16}$ notes. Furthermore, only monophonic melodies were considered i.e. one note played at a time. Another necessary constraint was to limit the pitches to the *C major scale* with 8 possible pitches so as to include one octave.

Each experiment set consisted of 30 tests of 500 generations each. Ten Zipf metrics (see section 2) were applied in the fitness function. The target vector is experiment de-

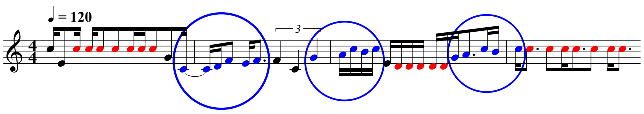

Figure 1: Example generated melody.

pendant and is thus described under the individual experiments. Other evolutionary parameters include population size (100), mutation rate (0.1) and crossover rate (0.9). The selection mechanism for both survival and reproduction was tournament selection with a tournament size of $k = 5$ and a probability $e = 0.1$ of selecting a random individual instead of the most fit.

3.1 Equal Metrics

As a first step, the target slopes of the metrics were all set to -1.0 (Zipf's ideal). The fitness of the best individuals at the end of each run ranged from 0.77 to 0.95.

Figure 1 provides an example evolved melody with the highest fitness (0.95) of the 30 runs. As illustrated, the melody includes long repetitive sequences of single notes. Examining all the results, such long repetitive sequences was seen to be a recurring theme. The fitness function seemed to be promoting such sequences, resulting in very monotone and quite unpleasant melodic progressions. Furthermore, many melodies that were subjectively seen as pleasant had in fact low fitness scores.

One possible cause of this feature lies in the metrics based on melodic intervals (including bigrams and trigrams). Sequences like these will increase the fitness associated with melodic intervals. However, these metrics also seem to support the emergence of melodic *themes* which help produce more coherent melodies. An example of a theme may be seen circled in Figure 1 where the sequence of intervals ($2 \rightarrow 3 \rightarrow 1 \rightarrow 1$) occur three times at different locations and pitches.

3.2 Realistic slopes

The first set of experiments applied uniform target slopes of -1.0. However, the melodies exhibited unpleasant long repetitions of notes. It was suspected that the cause of this monotony was the optimization of some of the Zipf metrics.

From research on Zipf metrics, one can expect varied slopes across the different metrics for a given piece of music. For instance, in the melody of Mozart's *Rondo alla Turca*, a range of metric slopes from -2.27 to -0.98 may be seen. To investigate whether such a given set of slopes might improve the generated music, the extracted slopes from Mozart's melody were applied, replacing the uniform target slopes from the previous experiment.

The experiments resulted in best fitness ranging from 0.52 to 0.95. Again, a subjective evaluation revealed even more monotony present in this set of melodies. The repeated note sequences of experiment set 3.1 was even more apparent in these results. This was likely caused by the generally steeper target slopes, which correspond to more monotonous distributions - see section 2.

3.3 Global vs Average Slopes

Mozart's piece is over 30 times longer than the generated melodies. Thus the slopes applied in experiment 3.1 and 3.2 can be seen as global characteristics of the music, which might differ from a 4 bar section of the same music. To investigate this issue further, Mozart's piece was divided into 32 parts, each of approximately 4 bar length and analysed individually for metric slopes. *Large variations in slopes were observed for each metric, with considerable deviations from the global metrics.* These fractal slopes were generally much flatter compared to their global variants, suggesting more randomness at a local level.

These results imply that such global metrics are very knowledge poor, providing little insight into music at a more local (smaller section) level. As such, global metrics provide little guidance for evolution, especially for generation of shorter melodies.

The mean fractal slopes for each metric, extracted from Mozart's melody, were then applied as the target slopes. An improvement in fitness was seen where fitness ranged from 0.88 to 0.99, suggesting that the flatter (more random) average fractal target slopes posed an easier task for evolution. The repetitive note sequences were less prevalent, which may be seen as a favourable property for avoidance of monotonous melodies.

4. CONCLUSIONS

The results of the experiments show that favourable musical features may emerge through the optimisation of several Zipf metrics. Further work is needed to refine the application of this technique for the generation of music. Although nice melodies can be generated, it is still challenging to generate pleasant music through the application of such a technique. It is thus believed that Zipf metrics alone may be too knowledge-poor to provide a workable fitness function for music composition. Further work will not only focus on further refinement of the technique but look at ways to incorporate more knowledge into the process, either through extension of the fitness function and/or through the representation.

Acknowledgements

The authors wish to thank the Classical Archives (`www.classicalarchives.com`) and Bill Manaris for their support.

5. REFERENCES

[1] M. Allan and C. K. I. Williams. *Harmonising chorales by probabilistic inference*. MIT Press, 2005.

[2] J. A. Biles. GenJam : A Genetic Algorithm for Generating Jazz Solos. 1994.

[3] B. Manaris, P. Roos, P. Machado, D. Krehbiel, L. Pellicoro, and J. Romero. A Corpus-Based Hybrid Approach to Music Analysis and Composition. In *Proceedings of the 22nd national conference on Artificial intelligence*. AAAI Press, 2007.

[4] G. Papadopoulos and G. Wiggins. A genetic algorithm for the generation of jazz melodies, 1998.

[5] G. K. Zipf. *Human Behaviour and the Principle of Least Effort: An Introduction to Human Ecology.* Addison-Wesley, 1949.