Generating Music With An Optimization Algorithm Using A Markov Based Objective Function

D. Herremans[1], S. Weisser[2], K. Sörensen[1], and D. Conklin[3]

[1]ANT/OR, University of Antwerp Operations Research Group, Antwerp, Belgium, {dorien.herremans, kenneth.sorensen}@uantwerpen.be
[2]Université Libre de Bruxelles, Brussels, Belgium, stephanie.Weisser@ulb.ac.be
[3]University of the Basque Country UPV/EHU, San Sebastián, Spain and IKERBASQUE: Basque Foundation for Science, Bilbao, Spain, darrell.conklin@uc.ac.edu

Abstract

In this paper, a machine learning approach is combined with an optimization algorithm to generate structured music for the bagana, an Ethiopian lyre. A low order Markov model that captures the style characteristics was built from a corpus of bagana music. Different ways are then explored in which a Markov model can be used to construct quality metrics that represent how well a fragment fits the chosen style. These metrics are implemented in the objective function of a variable neighbourhood search algorithm that generates bagana music.

Keywords: VNS, Music, Markov models, variable neighbourhood search

This research examines different ways in which low order Markov models can be used to build quality assessment metrics for an optimization algorithm that generates music. The metrics are used as the objective function of the algorithm which evaluates how good a solution or musical fragment fits the chosen style. A first order Markov model is built from a corpus of bagana music, a traditional lyre from Ethiopia, in order to capture the style of this type of music. Different metrics are then implemented in a variable neighbourhood search algorithm (VNS) that generates bagana music. It is often only possible to get rich statistics for low order models due to the size of many datasets. Unfortunately, these models do not handle structure very well, and often produce very repetitive output. Therefore, we propose a method to enforce structure and repetition within the generated music, thus handling long-term dependencies with a first order model.

We chose to work with music for bagana because cycles and repetition play an important role. In a first part of this research a method is described for representing and efficiently evaluating this structure with long-range dependencies provided by an existing template piece, while still employing a Markov model to evaluate the music.

In previous research, the authors generated counterpoint with a VNS algorithm using an objective function based on rules from music theory [5]. For most styles of music however, formal rules are not available. An overview of the different metrics implemented in this research, based on a machine learned model is given below.

High probability sequences Cross-entropy is used as a measure for high probability sequences, whereby minimal cross-entropy corresponds to a maximum likelihood sequence according to the Markov model [4, 6].

Minimal distance between transition matrix (TM) of model and solution Based on Davismoon and Eccles [3], an evaluation metric was evaluated that tries to match the transition matrices of both the original model and the newly generated piece by minimizing the Euclidean distance between them. This will ensure that they have an equal distribution of probabilities after each possible note.

Delta cross-entropy In order to optimize towards a cross-entropy value that actually occurs within the corpus, this metric optimizes towards the average cross-entropy value of the dataset.

Information contour is an additional constraint that describes the movement of information between successive events. High probability sequences often sound uninteresting and repetitive [6], this metric corrects this by enforcing diversity within a piece.

Unwords are significantly rare patterns [1]. They are the shortest sequence of notes (i.e., not contained within a longer unword) that never occur in the corpus, yet whose absence from the corpus is surprising given their letter statistics [2].

The metrics above are implemented in the objective function of a VNS that generates bagana music. The results are very promising and clearly indicate that high-probability sequences on their own are not sufficient. Yet combined with extra constraints such as information contour, a bagana expert described the generated music as "very good". The musical fragments generated with the TM distance metric were also among the better ones according to the bagana expert.

References

[1] Conklin, D., 2013. Antipattern discovery in folk tunes. Journal of New Music Research 42 (2), 161–169.

[2] Conklin, D., Weisser, S., 2014. Antipattern discovery in Ethiopian bagana songs. In: Proceedings of 17^{th} International Conference on Discovery Science, October 8-10. Bled, Slovenia.

[3] Davismoon, S., Eccles, J., 2010. Combining musical constraints with Markov transition probabilities to improve the generation of creative musical structures. In: Applications of Evolutionary Computation. Springer, pp. 361–370.

[4] Farbood, M., Schoner, B., 2001. Analysis and synthesis of Palestrina-style counterpoint using Markov chains. In: Proceedings of the International Computer Music Conference. pp. 471–474.

[5] Herremans, D., Sörensen, K., 2013. Composing fifth species counterpoint music with a variable neighborhood search algorithm. Expert Systems with Applications 40 (16), 6427–6437.

[6] Lo, M. Y., Lucas, S. M., 2006. Evolving musical sequences with n-gram based trainable fitness functions. In: Evolutionary Computation, 2006. CEC 2006. IEEE Congress on. IEEE, pp. 601–608.