

Automated Music Generation

Ryan Curtin

The Orange Lunchbox Brigade
ECE4884-L04: Prof. David V. Anderson

February 8, 2008

1 Introduction

The components of a music synthesis system range from instrument synthesizers to collections of audio samples to melodic generation. Recently, advances in computing processing power have made complex melodic generation and, more specifically, generation of entire works of music in real-time feasible. Many academic papers, articles, and dissertations have been written, describing possible methods of implementation for components of such a system and applications of music generation. However, commercial offerings for automatic music generation systems are very limited. This paper seeks to explore existing commercial implementations of automatic music generation systems and theoretical advances in that field.

2 Existing Commercial Developments and Patents Related to Music Generation

Perhaps the beginnings of music generation systems can be found in a system called iMUSE (**i**nteractive **m**usic **s**treaming **e**ngine), created by LucasArts developers Michael Land and Peter McConnell in the early 1990s [1]. iMUSE was developed for synchronizing music with visual action in a videogame; for example, if the protagonist of a game enters a car and starts traveling very fast, the music changes to become faster. While this system does not generate melodic ideas on its own, it is able to change between themes – an important capability of music generation. Later, in 1999, Cameron Browne patented a system and method for automatic music generation using a neural network [2]. According to the patent, the system is able to generate music based on an initial musical sequence given as input. In 2001, an unrelated free automatic music synthesis system called “AutoGam” was developed [3]. This system was capable of creating MIDI music given basic user input as to how the music should flow. However, its functionality is severely limited, as it appears to be an unfinished

abandoned project. As recently as one week ago, a patent was issued to Maryland inventor James W. Weider, who invented a music generation method, such that “each time a composition is played back, a different sound sequence is generated in the manner previously defined by the artist” [4]. While there have been several patents on subjects relating to automatic music generation, and a few commercially developed semi-automatic music generation systems, there are no completely automatic music generation systems currently on the market.

3 Theoretical Advances in Automatic Music Generation

In the academic world, several papers have been written concerning automatic music generation. These cover topics ranging from fully automatic generation of music based on existing music to modeling musical styles with methods derived from machine learning. For example, Dr. Ulf Berggren of Uppsala Universitat published a paper describing a method for algorithmic construction of sonata movements [5]. This method uses Mozart’s piano sonatas as a model for new sonatas; however, the system is restricted to classical music generation. A more general method of sequence design, applicable to such things as NMR data interpretation for protein structure determination, modeling Internet traffic, and melody generation, was designed by Dr. Manan Sanghi of Northwestern University in 2006 [6]. Machine-learning methods were used as a basis for musical modeling in an article produced in the IEEE Computer Journal [7]. The writers suggested that this could have applications in computer-aided music composition, musical generation, and musical prediction. A more qualitative study of the components required in a music generation system was undertaken by Scott Downie, at the University of Kansas [8]. Downie specifically studied a generation method to create a “soundtrack” for a customer’s personality, and documented the factors that needed to be considered when creating such a system.

4 Design Methods for Music Generation Systems

Since music generation is still an emerging field, many methods exist for designing a music generation system. Probability and statistical modeling, though, must be present in every system if it is to produce music with any degree of randomness. Dr. Ulf Berggren’s method of sonata design used a probabilistic model based on Mozart’s sonatas; this model was generated by analyzing a library of Mozart’s sonatas. These probabilistic models may use concepts like Markov chains and Hidden Markov Models (HMMs), which are normally found in voice recognition applications [9, 10], to analyze music and generate probabilistic models. Another method of music generation, patented by Cameron Browne [2], used a neural network for music modeling. In this application of a neural network, the ‘artificial neurons’ are interconnected and use a complex computational model to process

input (in Browne’s patent, this was an initial input sequence of music). It should be noted that as a music generation system becomes more flexible, able to produce a variety of types of music, and able to produce technically complicated and artistic music, the probability models driving the system will become incredibly complex and large, making it difficult to keep the system producing music in real-time.

5 Components of a Music Generation System

Implementing a music generation system is most likely to be done in software. While it could feasibly be done on a number of different FPGAs and ASICs, they present no real advantages to music generation. Pre-existing libraries for music analysis, probability modeling, and other mathematical functions that would be required in a music generation system could be utilized if the system was written in software. Therefore, a multipurpose CPU presents the best platform for developing such a system. For most of the patented music generation systems, implementation details are limited at best, and their choices of libraries, language, and platform for implementation are not given. Berggren’s Mozart-based sonata modeling system was written in software, using the Prolog programming language. Many other automatic music generation systems proposed in conference proceedings and research publications do not offer implementation details but only conceptual designs and formulas.

References

- [1] “iMuse Island: What’s iMuse?.” <http://imuse.mixnmojo.com/what.shtml>, 2004.
- [2] C. Browne, “System and method for automatic music generation using a neural network architecture.” U.S. Patent 6,297,439, Aug 1999.
- [3] “Autogam.” http://autogam.free.fr/a_titre.htm, 2001.
- [4] J. W. Wieder, “Generating music and sound that varies from playback to playback.” U.S. Patent 7,319,185, Jan 2008.
- [5] U. Berggren, *Ars combinatoria: Algorithmic construction of sonata movements by means of building blocks derived from W. A. Mozart’s piano sonatas*. PhD thesis, Uppsala Universitat, 1995.
- [6] M. Sanghi, *Sequence design and discovery*. PhD thesis, Northwestern University, 2006.

- [7] S. Dubnov, G. Assayag, O. Lartillot, and G. Bejerano, “Using Machine-Learning Methods for Musical Style Modeling,” *IEEE Computer Journal*, vol. 36, pp. 73–80, Oct 2003.
- [8] S. Downie, “Motion picture moods for geeks and non-musicians: An automated, intelligent music-generation system,” Master’s thesis, The University of Kansas, 2006.
- [9] J. Dai, J. Tyler, and I. MacKenzie, “Application of Markov chains to speech recognition,” *IEEE Electronics Letters*, vol. 27, pp. 2360–2361, Dec 1991.
- [10] R. D. Mori and F. Brugnara, *Survey of the State of the Art in Human Language Technology*. Cambridge University Press, 1997.