Reid Swanson
ISE 575

AI Methods for Algorithmic Composition

Papadopoulos and Wiggins give a survey of AI techniques that have been used in trying to generate music. They also provide some analysis of the previous work by examining various shortcomings, considerations and possibilities for future work.

They split the techniques into five separate areas of research. The first that they examine are methods based on mathematical models, which include Markov chains, chaotic non linear systems and some machine learning techniques. Although these models are quite popular because of their ease and complexity, especially Markov models, they suffer from a number of drawbacks. Acquiring the probabilities, deviating from the norm and representing higher level knowledge are all difficult tasks.

Knowledge based systems are the next approach they consider. They are primarily categorized by their use of symbolic rules or constraints to solve the problem. The major advantage they have is their ability to explain the decisions made in generating the piece. However, engineering the knowledge is difficult, time consuming and depends on th expert.

Grammars are another way of algorithmically composing music, where one tries to specify the regular structures inherent in the music and exploiting them to analyze or make new music similar to the way some language research is done. Unfortunately, their inherent hierarchical structure limits their ability to “improvise” effectively. Additionally parsing is computationally expensive and it is hard to distinguish the many “bad” possible strings that can be produced from the good ones. Although the paper makes a claim about a lack of semantics involved with grammatical approaches, I don't see how it is any worse than the other approaches. In fact the line between syntax and semantics is often hard to define since the semantics of something are often highly dependent on the syntax so it would seem that grammatical models are actually offering quite a bit in terms of semantics, although perhaps not explicitly.

Evolutionary methods are mostly based on genetic algorithms and are split into two categories. The first are based on objective fitness functions, while the other require a human to provide the fitness of a given generation. While the first type require less intervention, and so less cost, time and effort, they do not always capture the human process and lead to strange results. On the other hand human intervention requires much more effort and is very subjective.

Machine learning and artificial neural networks provide the framework for systems that learn. They tend to be good at learning shallow features but fail to generalize or capture higher level knowledge. Finally they show several techniques that try to leverage the strengths of several different approaches. Most notably with ANNs and constraint satisfaction.

There are several issues that the paper address that are important to all the methods described but are not necessarily addressed by the individual systems implemented. They first note that for nearly all the systems covered there is a lack of a consistent, formal or scientific evaluation methodology. We

are essentially left to their word that it works or doesn't, and do not have any consistent measure by which we could compare across systems. Although I agree this is fundamentally important, since we are only stabbing in the dark without, I also believe it is one of the hardest things to do. Not only is it expensive and time consuming but the quality of music is highly subjective. Getting a panel of experts to agree would be quite challenging unless very specific qualities were identified to judge upon (e.g. rate the fluidity of the melody, or consistency of the rhtym). However, the more specific these conditions become the less the process becomes about generating “music” as it does about satisfying a constraint. For example, what makes a good melody does not necessarily make good or interesting music.

It is also important to consider how the data is represented. Not only is it important to consider for efficiency as some of the work did but it is also important that some tasks may be better suited with different representations so it should be more beneficial for a system to be able to dynamically change it's representation to best fit the goal. Similarly they not that goals are important and that most attempts have been framed as means of problem solving, but that it might be more important to frame systems in terms of a creative process containing goals and intension.