Task 16 - Analysis of Music Samples and Collaboration/User Friendliness

This task provides my thoughts on the final outcome of the project.

On User Hostility

As the sole user of this system, I can confidently say that it is not user-friendly. I spent five hours generating the seven songs of the “Synthetic Soulscape” album in Penfield library on a Friday (that’s dedication, or perhaps a sign of insanity). The interface is clunky, especially with its reliance on ABC notation. While MuseScore was a definite step up from EasyABC, its ABC import plugin created complexity by insisting on opening a new window for every new music input. I think things could be improved if everything was kept directly within the app, as opposed to relying on a third-party to play samples. This could be further improved with a nice GUI.

Genetic Algorithms are Difficult to Fine-Tune

With this project especially, it is difficult to navigate the line between variation and convergence. Variation is to keep user interest and introduce samples that a user may like over that of the initial population. However, too much variation detracts from the point of the genetic algorithm, and we lose the user's influence on the musical outcome. On the other hand, too much convergence results in an incredibly boring experience for the user and – in the same way – removes the user’s influence on the musical outcome.

I was unable to crack the best parameters for striking this balance. I did, however, make the following observations:

1. If your population is too small, convergence occurs quickly.

2. If the music samples are too long, it is difficult for the user to deduce whether or not a sample is a mutated sample or a copy. I think I got tricked by this a few times and incorrectly assumed that a sample had converged, when it hadn't.
3. It is difficult to breed out the longer melodies resulting from mutation in future generations.
4. Increasing the number of generations a user has to rank is quite exhausting for the user. My longest listening session was over twenty minutes for ten generations.
5. Increasing the amount of rest a user gets (i.e. ranking every couple of generations) makes it feel like there is no user influence on the outcome.

I am not sure if there are feasible grounds for collaboration with algorithmic composition using the hybridized genetic algorithm and constraint system I created for this project. I do have some ideas on how to build on this design, which is discussed in the final section of this paper.

Is Collaboration Actually Happening?

The latter observations beg the question: *is collaboration actually happening?*
My answer: I do not think there is enough data to truly tell.

A new question: i*s collaboration feasible for this proposed system?*
My answer: I do not think so. As I mentioned previously, the parameters for the genetic algorithm are very difficult to fine-tune for this specific problem. If there is a way to achieve collaboration, it may be too taxing on the human, like involving an insane number of rankings or something along those lines.

Now, what?

Final Thoughts

I think there is some merit in using genetic algorithms in conjunction with other algorithmic composition systems to increase user collaboration/influence on musical outputs. The initial paper I read mentioned using Markov chains, rather than a constraint system using this approach. I am not sure if this approach would help the aforementioned issues I had.

However, I do think more interesting results could be achieved using other approaches.

I have been doing some research on Hidden Markov Models and algorithmic composition, and have played with a program that trains on a composition and outputs a similar composition in the style of whatever was inputted. The curious part of the program is that it outputted results that sounded different from the original, yet similar enough that you could connect the dots if told what the original input was.

Perhaps, using an HMM would be a better mechanism for mutation, as a way to increase variation while still maintaining the style whatever was chosen. The HMM I looked at had some metrics for measuring similarity to the original piece as well as musicality, so perhaps those could be employed in some way as a fitness metric for the genetic algorithm. None of this is fleshed out, obviously, but I think there are grounds for experimentation in this area.

All in all, I have not given up hope on the potential for music collaboration between artificial intelligence and humans.