Adaptive Elusion - an improvisation for pianist and real-time machine learning

PALLE DAHLSTEDT, University of Gothenburg and Chalmers, Gothenburg, Sweden

1. PROGRAM NOTES

A continuation of my experiments with minimal algorithms, investigating how small an interactive musical algorithm can be and still invoke the feeling of "somebody there". Here, a small set of adaptive algorithms react to a live pianist, trying to imitate, elude and counteract his playing, while at the same time being completely dependent on it as a source of patterns and sounds. The piece explores real-time training as a primary modus of interaction, in a cat-and-mouse game of sorts. It is also an example of what I call entangled musicianship. What the pianist plays is a reaction to what the algorithm plays, and at the same time shapes the future playing of the algorithm, hence entangling performance and control.

The musical response is generated by a small machine-learning algorithm that starts empty and is trained in real-time on what I am playing. It can also gradually forget what it has learnt.

Fig. 1. Image from the premiere of Adaptive Elusion, at the Royal College of Music (KMH), Stockholm.

2. PROJECT DESCRIPTION

This interactive musical work is part of a larger research project into interaction with *minimal* algorithms, i.e., lean interactive algorithms that get their "liveness" and perceived agency from the performer, but it is transformed enough so that it feels like there is somebody else there.

By experimenting with different kinds of algorithms with such properties, we explore the conditions of agency, both as sensed agency for the performer and perceived agency from the algorithm, evaluated in relation to interaction and algorithm design.

In Adaptive Elusion, a small set of machine learning (ML) models start completely empty, and are trained by everything the performer plays. Each phrase played is 1) heard in the room and 2) used as training data, altering the future behavior of the system, with binding consequences for its immediate future behavior, and finally, it 3) triggers new phrases to be generated by the algorithm. I call this situation *entangled musicianship*, where each action has multiple meanings and complex future implications. Together, the human and the algorithm create a time-braid of actions, reactions, and re-reactions, in a situation where I cannot avoid the responses and gradually changing behavior of the algorithm.

In the algorithm, forgetting is also implemented. With a certain probability (adjustable by the pianist through his playing), the model "consumes" its own capabilities, which gives space for new learning, and keeps the interaction between human and algorithm continuously renewed.

Throughout the performance the pianist explores the affordances of the system, and tries to navigate the system, in all its unpredictability, in the search of interesting and meaningful music. This search in itself becomes a kind of narrative form, shaping the musical form.

As the models are small, each musical interaction has a perceivable impact on the behavior of the algorithm through immediate training, and in contrast to today's typical huge AI models, the lack of pre-training creates a more intimate relation to the model, emphasizing direct interaction, enhancing the sensed agency for the performer, while maintaining a perceived agency from the algorithm, in spite of its small size and limited complexity. This works through the capturing of vital liveness qualities from the performer, and mirroring these back in a transformed state, without losing the musical qualities.

So, the piece is an experiment in how to design an algorithmic interface for musical co-expression, where the performer and the algorithm form a system together, in which the complexity emerges from the carefully designed entanglement between the two.

The system is implemented in Max, with a set of simple hardcoded custom ML models (of Markov variety). This particular algorithm was selected to have access both to direct training and forgetting of specific learned behaviors.

3. TECHNICAL NOTES

The piece is a live performance involving one pianist (the author), a grand piano, and electronics. Ideally, it should be performed on a Disklavier player piano, as the algorithm response will then be performed on a real piano, in the same medium as the live action. However, in the lack of a player piano, the premiere (in the linked video) was performed on a grand piano using a Moog Piano Bar MIDI sensor (designed for concert grands, placed on the keyboard near the lid), with the algorithm running on a laptop next to the pianist, and algorithm performance realized with a grand piano plugin. This is a perfectly reasonable configuration. A performance on one normal grand piano and a Disklavier player piano (for the algorithm response) could also be considered, but is by no means required.

Duration: 10-12 minutes

I will bring:

- Pianist (myself)
- Moog Piano Bar MIDI sensor (if you don't have one)
- laptop with audio interface, 2ch on unbalanced 1/4" jacks

From NIME I need the following:

- a tuned grand piano, 88 keys
  (optional: or a Yamaha Disklavier grand piano if available, or both)
- 2ch speaker system
- amplification of the grand piano (if needed to achieve a good balance with the generated piano part, depending on the acoustics)
- monitor speakers

4. MEDIA LINK(S)

Video link: https://youtu.be/kclJz0j1cto?si=7Yw_-1b1zkPHSBZA

ETHICAL STANDARDS

This work involves no other humans than the author, and the artistic research funding behind it was obtained under standard ethical regulations of the Swedish Research Council. The AI/ML algorithms used are too small to be sentient, and no algorithms were harmed in the process.