We have made substantial progress on the project we proposed for the Agile Seed Funding. Our system uses Python with MediaPipe to detect hand landmarks (specifically finger spatial position) in real time and transmits this data as OSC (Open Sound Control) messages to custom Max for Live patches running inside Ableton Live. We were able to recruit two summer undergraduate student interns who came with their own substantial stipends: Deepansh Goel via the Google Summer of Code, and Alison Wang via a MITACS Globalink Award. We were therefore able to use the Agile funding, along with some supplemental funding from our own resources, to get better versions of the hardware we had originally planned for: a high-quality ultra-wide varifocal lens (for the high frame-rate camera that we already had) and a laptop with an Nvidia RTX 5080 GPU that will enable the portable use of our system with fewer limitations. Deepansh built a full Teensy-based setup to measure the gesture to sound latency and benchmark the system, and we could achieve a median latency of around 15 ms, and also substantially improved the efficiency of the camera capture software as well as the gesture to OSC command translation using multiprocessing techniques. Building on this, Alison created several prototype components as well as interfaces to map the OSC commands to music signals. Alison wrote a Max for Live patch that converted the incoming hand coordinates into virtual dials that musicians can map to any audio parameter, filters, effects, volume, and reverb. Alison then adapted FluCoMa's neural network synthesis controller to enable complex mapping of hand positions to 10-parameter chaotic synthesisers. She also built an OSC-based gesture recorder which listens to incoming OSC messages on UDP port, captures the complete gesture stream with timing information, stores the data, and allows playback at adjustable speeds or backwards with delays. Alison also worked on a simple spatial drum mapping patch that implements quadrant-based drum triggering for air-drum performance control. She also explored corpus-based sound browsing by developing a neural corpus synthesis instrument using MLP Regressors from SPTools, which enables real-time sound selection and playback through hand gesture navigation across a 2D corpus map. Alison worked on a four-speaker spatial audio system that enables musicians to move sound around a room using finger-pointing gestures, and modified Eric Ameres' Max for Live quadraphonic panner to accept OSC input from gesture recognition, implementing an equal-power panning algorithm that maintains constant perceived loudness across all four speakers regardless of position. These are examples and proof-of-concept prototypes that we will use to interface with musicians who wish to work with our system. We have already spoken with leading thereminists in France and Montreal who are very excited about the possibility of a digital theremin: we aim to set up a collaboration to create mappings, possibly learned with AI methods trained on bespoke videos of theremin performances, that will enable a gesture-controlled digital theremin that may get around the capriciousness and finicky nature of real physical theremins. We are also in touch with the music program at UdeM about the possibility of their students using our tool for their course projects. We are also currently preparing a manuscript detailing the device and its latency characteristics, and we plan to conduct experiments on the impact of different gesture-sound delays on the feeling of agency and naturalness.