Did You Hear Something?

Exploring and modeling perception, cognition and ability through the installation *Why am I doing this to myself?*

Pat McMaster

EAST 407, Assignment C

March, 2022

After the Silence (Études)

After working on the Silence Études in assignment B, I felt it was time to move into new territories and answer some of the questions that came up during my research and from members of the class after my presentation. Questions that particularly stuck with me were:

A. How would these experiments work in a noisier environment?
B. Could the piece be performed in a way that responded dynamically to this environment?
C. Would listening-related social pressures change if the experiments were conducted in an installation format vs a classroom performance?
D. Would the ability to perceive quiet sounds change if human voices were included?

To answer these I would need to gather resources and to consider ways in which to build off my previous experiments this semester. What came to me at the beginning of this process was the idea of developing a feedback system which would listen to the ambient sound of the room and dynamically control the relative volume of the piece as it is performed in the space. I sketched this out while riding the Métro home from class, and created a proof-of-concept in VCV Rack using a recorded crowd-noise sound and one of the Silence Études. I'll go into more detail on this process in a later section, but suffice it to say that small success early on was crucial to emboldening me to move further outside the familiar methodologies of my practice and into planning an installation. Once I knew my presentation date, I immediately booked a practice room at Concordia to host the installation. This meant that I would have an anchor firmly grounding this decision. I booked speakers, microphones, stands, projectors and screens in sufficient numbers to allow me some flexibility as I tried to answer the big questions ahead of me: what materials and resources would need to be gathered, and what themes would I want to explore / what story would I want to tell through this installation?

The mind-body problem

The note about embodied cognition mentioned in the comments for my Assignment B submission haunted me for a good while. Unfamiliar with the term, I found some resources to help explain the continuum of philosophical thought contained within those words. I found myself again on the uncertain ground of radically new ways of thinking, experiencing and relating to the world around us. I am always disappointed at the defensive pattern I seem to find myself in when this happens. While I'm sure it serves some evolutionary purpose, as I get older I have to ask: was it always like this, did I always display a kind of resistance to having my perception challenged? I found myself bristling as I went down the embodied cognition rabbit-hole more than I care to admit, but ultimately this friction is a major part of why I am so glad to be a part of EAST 407 and Concordia in general. I came to university to grow, to be stretched, to encounter new ideas and discover new ways of seeing things, and this is inherently a messy, disruptive, uncomfortable process.

After my initial resistance, I sheepishly discovered that rather than having my flag firmly planted in the camp opposite embedded cognition, I had in truth very little connection to the Cartesian mind-body dualism which sees the body, including the brain, as being separate from our consciousness and mind. The classical idea that the brain was a kind of organic computer, processing the stimuli derived from the senses in a given environment and running a fixed "operating system" which controlled our actions was also one that I had always struggled to fully accept. While embedded cognition covers an absolutely staggering number of ideas of which some felt especially alien, a great many of them aligned with how I'd imagined the relationship between mind, body and the environment. A summary that resonated particularly strongly with me was that "many features of cognition are embodied in that they are deeply dependent upon characteristics of the physical body of an agent, such that the agent's beyond-the-brain body plays a significant causal role, or a physically constitutive role, in that agent's cognitive processing."[1]

I have been wrestling my own internal demons concerning the relationship between brain, body and cognition since I first became ill at the end of 2015. Intense migraines with crippling prodrome, chronic pain in my lower body, sleep disturbances and eventually cognitive disabilities followed. I experienced overwhelming tiredness and had immense difficulty focusing on my work, remembering names, details or words. I went from being a model employee working at a high level of information technology specialization to having to check my work three or four times before submitting it due to cognitive

[1] Wilson RA, Foglia L (2011). "Embodied Cognition". *The Stanford Encyclopedia of Philosophy.*

errors. After multiple failed attempts at accommodation with my employer, these difficulties culminated with my taking a long-term disability leave in February of 2019 and focusing full time on finding out what was wrong and how to treat it. The near-constant pain I experience in my lower limbs and extremities has no basis in the muscles, joints or bones themselves. There is a disconnect between the body's state and the mind's interpretation of that state. My body never seems to be able to get enough sleep, but also struggles with bursts of disruptive insomnia. I have cataplexy, a rare disorder that is almost never found outside of a narcolepsy diagnosis, and yet after three overnight sleep studies it was concluded that I do not have narcolepsy. I have periodic limb movement disorder, in which - unbeknownst to me - my legs have been moving on their own accord while I'm asleep for goodness knows how long. Most recently I've developed propriospinal myoclonus, a rare disorder which produces arrhythmic jerks in my trunk or abdomen. For all of these symptoms and others, there is no formal diagnosis nor treatment. A doctor told me that fibromyalgia could be considered a 'diagnosis by exclusion' because nothing else could be found wrong with the parts of my body where I experience disruption. The sleep specialist doctor stated 'idiopathic hypersomnia' as the catch-all for my puzzling sleep activity. The interconnectedness of the brain, the body and the mind remain elusive.

The concrete model of a healthy brain transmitting and receiving electrical impulses through a nervous system stretched throughout the body which I grew up believing has been forever challenged by these recent changes in my abilities and onset of chronic pain. What organ or organs, what system or systems are working abnormally enough to cause all of this? So much of our body is designed to help us avoid pain, and yet my system is causing it seemingly without external influence. I struggled with this self identification, and often found myself wanting to ask "Why are you doing this to yourself?" My mind, acting as interlocutor, questioning the capital-b body's mysterious and self-defeating behaviour. I don't believe this mind-body separation to be truthful, or even terribly valuable even as a metaphor to my accepting my current situation, and so the question that I carried with me as I gathered resources and explored visual and sonic territories was "Why am I doing this to myself?" The mind capable of asking this question indivisible from the experiential corpus that through some unknown mechanism or flaw is actively found to be in an oppositional, adversarial position with itself.

I have found comfort in wrestling with these questions in the members of my family, some of whom are also living with a fibromyalgia diagnosis, and who have similarly struggled for decades trying to find the right specialists to listen to their stories and administer the right tests and treatment. I have no religious faith, and see no particular virtue in my suffering, but I try to express gratitude for the paths it has opened for me, chiefly my ability to focus part-time on an academic experience that had previously been shut to me.

I've been working with machine learning and artificial intelligence in a creative discipline for a number of years now, and am fascinated by the human meta-narrative that emerges from the analysis of trillions of data points used to train algorithms to understand how we think and behave. I have asked medical experts and specialists why the symptoms I experience have increasingly taken hold of my day to day life, and there have been very few answers. In the various support groups and online forums devoted to my disorders in which I've spent time, no clear pattern or satisfying resolution seems to exist for any of us. The desire for answers seems to have become too much for some of the people I've encountered, who have been led astray by snake-oil salesmen and con artists offering miracle cures where none truly exists. What perspective could be gleaned through the lens of machine learning? How would these tools reflect my lived experiences back at me, and would I be able to find some connection to what these non-human tools offered me?

As in my other AI projects, I considered these processes to be a collaboration between myself, the software and the audience. This helped me feel more supported as I explored the more vulnerable elements of the resource gathering that followed.

It always takes a certain amount of time to get back into the machine learning headspace, and so I eased in by generating the text that would be fed to the text-to-speech voice generation software. The first two generated texts were strings of words strung into nonsensical sentences, but containing English-friendly rhythm and punctuation. The other three were more advanced - generated texts that read like excerpts from short stories and were decidedly prose-like. I provided these five texts to different voice engines in order to provide some variation in gender and age for the resulting sounds. The listener's ear would have more variety of material to grab hold of that way, and I found it interesting to see how the timbres of certain voices felt more 'real' than others, despite being of similar technological pedigree.

In case the human voice gambit proved fruitless, I wanted to shore up my composition with one last exploration of 19-equal division of octave microtonal music. I had yet to explore chord progressions within this temperament, and so developed a patch in VCV rack that would modulate between different chords and voices within this alien collection of pitches. I found myself returning to the more conventional intervals and chords, but managed to push myself into more experimental ratios of frequencies in this last kick-at-the-can with 19-edo.

I took the five speech elements and lightly processed them so that they could still be identified, but were recognizably transformed. I took these processed voices and the 19-edo microtonal composition and heavily worked them into a glitched barrage of sound that would serve as a third compositional element. I stretched this resulting sound until it was about 75 minutes long and ran it through a filter bank so that only the highest and lowest frequencies were inter-modulated, an homage to the extra-frequency range work of the earlier assignments. This would serve as the foundational bedrock for the installation, made up of the other component pieces but on a completely different timescale, so that there would be a sympathetic macro-length compositional element present at all times.

I then returned to my self-listening, proof of concept patch from earlier in the month, which I had nicknamed “Shy Guy.” In the flowchart below you can see the principles at work.

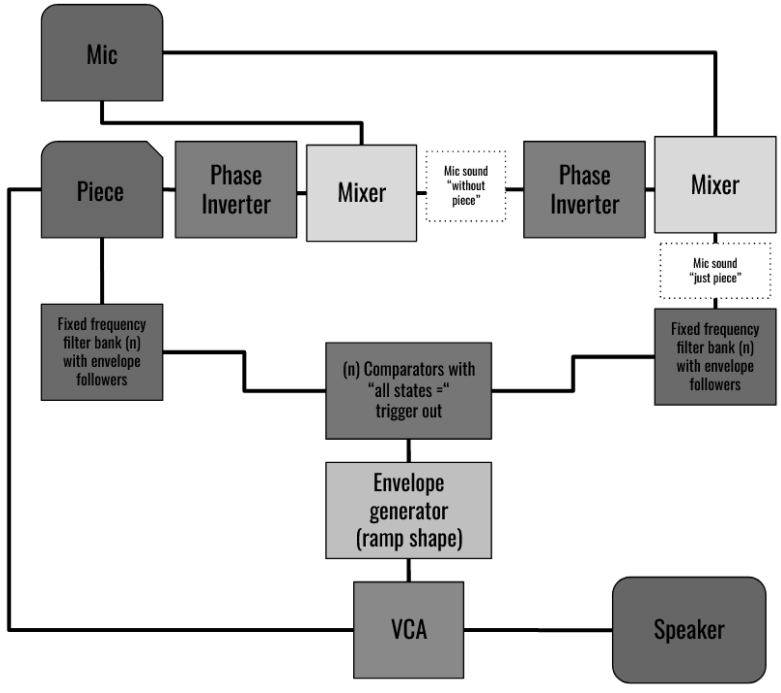

Through phase inversion of the live in-the-room sound, fixed frequency filter banks with envelope followers and a multichannel comparator, the patch “listens” to the room, compares what it hears to the

reference copy of the piece, and when it can ‘hear’ the piece out in the room, it lowers its volume to zero, and then slowly raises it again, repeating the process.

I tried it out with the completed compositional elements, and with some tweaking of volumes and sensitivity, it did what it was supposed to do! I would try four “Shy Guy” instances with a quadraphonic set-up and a single instance taking feedback from all four microphones in the workshopping of the installation at Concordia, but otherwise the sound aspect felt completed.

For the visual elements I wanted a combination of text and images based on prompts provided to machine learning algorithms. For the text I started with the prompt “Why am I doing this to myself?” and recorded about a dozen extracts ‘completed’ by the AI. They were all very interesting, but the five I selected can be found below.

Why am I doing this to myself? I've got to stop. I've got to stop this insanity, before I put myself or someone else in danger. I know I'm nowhere near ready to hunt a vampire, but a woman can dream, right?	Why am I doing this to myself? I'm not like this at home. If it's hot in the summer, I'll usually stay indoors. Why must I inflict this kind of torture on myself? The last thing I want is for my friends to start referring to me as "Mr. Chilly." No, thank you. Most days, I'd just rather be a hermit.	Why am I doing this to myself? I cannot stop myself from compulsively checking my phone. Sometimes I'm checking it in bed. What kind of person is that? I've been neglecting my reading. I have a stack of books, all of which I started to read and put down.	Why am I doing this to myself? The mountain top seemed very far away. I began to question my decision to climb that damn mountain. The top is totally hidden behind trees, fog and snow. It is icy and the path is hard to see. There is no one else to help pull you up the mountain. There are no other people to show you the way.	Why am I doing this to myself? I'm not one for easy explanations. But a few weeks ago I came across this particular explanation that speaks to me. I know I am not alone in this dilemma. I also know it would be easier to simply cut my losses and stop the hunt. But I've spent so much time...

In the first two, I appreciated the humour - I hesitate to call it ‘unintentional humour’ because I don’t pretend to understand the ‘intentions’ of the AI, if any. I like to think that an endearing sense of humour might emerge from our experiments in machine learning. The second group of two resonated strongly with me, the first for the banal self-admonishment concerning a relatable overreliance on technology - particularly striking to emerge from an extremely sophisticated piece of software! The second reminded me of when I was solo climbing Taranaki Maunga on the West Coast of New Zealand’s North Island, a

particularly important memory that seemed perfectly encapsulated in this quote, to say nothing of its resonance with the feeling of the metaphorical medical mountain I am climbing. The last extract I ultimately rejected because it fell too sharply into what I consider the AI uncanny valley. I have a document of such machine learning text generation results, ones that cut too close to the bone, too human-like in their discussion of important subjects, or express ideas that are too controversial to risk being mistaken for expressions from myself or my collaborators, and so I would prefer that these not be included in such a personal project.

My last challenge was the most rewarding - to see how a number of ideas ranging from the concrete to the very abstract would be represented by text-to-image generative adversarial network algorithms such as DALL-E Mini. I submitted the following prompts: nervous system, brain, neurons, hypersomnia and fibromyalgia and collected between 21 and 40 generated images per prompt which I chose for their aesthetic value. I've included the 20 selections from each prompt with this report, and am truly fascinated by the abstract and beautiful representations that were generated out of these sophisticated tools. The very imperfections in the way in which they represent the subject material are what makes the resulting artwork so interesting, and I tried to internalize that as I struggled with the subject matter of this project and its impact.

I created a looping video which showcased these images and the text extracts, and armed with that and my sound components I was ready to put everything together in the Concordia studio and welcome people to the installation.

Learning by doing

I had two hours to get the installation set-up, and fortunately Kevin Moon was extremely generous in helping me not only secure a better venue than the original crowded practice space I had booked, but also to help get the physical installation of speakers, microphones, projectors and screens up and running. This gave me the freedom to work with the space and the positioning of microphones and speakers, to tweak the patch and consider different spatialisation options, rather than just running cables for 90 minutes.

My laptop did not handle the four "Shy Guy" instance setup well. The CPU was maxed out and there were heavy audio buffering artifacts. Even when I decided to move the hour long bedrock track to the studio's computer and 5.1 setup, the quadruple live processing and playback was just too much. I moved to my backup plan, which was to sum the inputs of the four microphones and combine them in such a way that the live room sound was being compared to the combined voice tracks, and this worked much better. In addition, I found that having a single combined source of voices that moved throughout the space was easier to track spatially as a listener than four sources. The microphone sensitivity was the only part that was a bit difficult to prepare for in advance. I found that if I clapped my hands or produced multiband sounds at a loud enough volume - through the use of white noise for instance - that it was a close enough approximation for crowd noise. The challenge became what would happen if the room was too quiet.

In the absence of noise to compete against, as soon as the volume started to increase on the voice component, it would hear itself and lower the volume immediately, and so the biggest fear was that people would be "too polite" and that the voice sounds wouldn't ever enter the space. I wrote in a failsafe, where if the amplitude hadn't changed in 3 minutes it would trigger the cycle, and this was helpful when I was first trying things out without anyone else in the room.

When the first people arrived to check it out, I made it a point to engage in small talk so that they knew that as the composer was breaking the silence, they could feel free to make sound in the space. I experimented with simply letting one student in on their own while I wasn't in the room, and they did remain completely silent. Then two students were invited in without my circulating, and they both remained silent; it was easy to tell because I could listen for the telltale voice components while waiting outside. The non-silence nudge was necessary, and so for all the subsequent experiences in which I could be party, I made sure to talk with people in the space and allow their sound to build sufficient

volume that the “Shy Guy” patch could perform its magic. The amount of sustained volume that would be required to have the vocal element present for a longer period of time would really be quite a lot, and far beyond what I think even the noisiest single installation visitor could provide, and this was useful data to know for next time. I kept my lips sealed as to what was going on, because I was curious to see if anyone would notice or interact, but I fear that the sample size of students - there were many COVID-related absences that day - and the relative quiet of those in attendance meant that it didn’t appear that anyone noticed a relationship between the ambient volume of the space and the appearance of the voices.

The social convention programmed into the “Shy Guy” patch is only one possible way to behave of course - I considered a “Rude Dude” patch that would do the opposite, increase the volume to be heard over anyone talking, but I’ve worked with enough electric guitarists over the years so I don’t think the world really needed to have that behaviour modeled, even if it could be turned off at will!

While reflecting on how things went, I became curious about how incorporating visual sensors could add another layer of sensory embodiment to the project, and so mocked up an image of a three-screen/wall projection installation with Kinect cameras or similar infrared proximity detecting sensors that could affect the visuals that were displayed on the screens. I imagined a fuzzy appearance effect applied to the text unless the viewer moved closer to the wall, upon which it would become focused and sharp. It would be ideal to have the microphones and projectors located on the ceiling so they wouldn’t interfere with the movement throughout the space. These experiments would require some more powerful computers and resources, and hopefully one day I can get a chance to develop some of these ideas further.

The responses from those who experienced the installation were really encouraging and positive, and I felt like I learned a great deal about the importance of pre-site preparation and leaving enough time for the ‘unknowables’ that will present themselves in the room, imagining it as another collaborator. I would definitely like to work more in this medium, and feel like I’ve only scratched the surface on the relationship between machine learning, embodied cognition and my own personal relationship with the ideas of mind and body. A number of students reached out to me after the presentation with some great ideas and directions I might want to take this research-creation in the future, and so while this is the end of EAST 407 I know that the work I’ve done this semester will have a lasting impact on the direction I find my practice taking me.

References

Ceko M, Bushnell MC, Gracely RH. “Neurobiology underlying fibromyalgia symptoms.” *Pain Res Treat*. 2012;2012:585419. doi: 10.1155/2012/585419. Epub 2011 Oct 27.

Cox, Arnie. 2016. *Music and Embodied Cognition*. Indiana: Indiana University Press.

Thacker, Mick. 2015. “Louis Gifford – revolutionary: the Mature Organism Model, an embodied cognitive perspective of pain.” *In Touch*. 152. 4-9.

Sarzi-Puttini P, Rizzi M, Andreoli A, Panni B, Pecis M, Colombo S, Turiel M, Carrabba M, Sergi M. “Hypersomnolence in fibromyalgia syndrome”. *Clin Exp Rheumatol*. 2002 Jan-Feb;20(1):69-72.

Shapiro, Lawrence. 2014. *The Routledge Handbook of Embodied Cognition.* New York: Routledge.

Small, Christopher. 1998. *Musicking: The Meanings of Performing and Listening*. Hanover: Wesleyan / University Press of New England.

Tools used

DALL·E mini - Generate images from a text prompt

https://huggingface.co/spaces/flax-community/dalle-mini

IBM Watson Text-to-Speech - Convert text into natural-sounding speech

https://www.ibm.com/cloud/watson-text-to-speech

InferKit - State-of-the-art text generation

https://inferkit.com/

MeldaProduction - Professional audio software

https://www.meldaproduction.com/

VCV Rack - The Eurorack Simulator

https://vcvrack.com/