

## **AI/ML in the Sonic Arts: Pitfalls and Pathways**

*Stephen Roddy, Radical Humanities/Digital Humanities UCC. srodgy@ucc.ie*

*Robin Parmar, DMARC, CSIS, UL. robin.parmar@ul.ie*

This commentary considers the application of Artificial Intelligence (AI) and Machine Learning (ML) technologies to music and the sonic arts. It critiques the classical computational theory of mind (CCTM), a doctrine deriving from functionalism, which codifies “mind” as a mathematical function symbolic representations from one dimension (mind) can be directly mapped onto another (world) in accordance with a given transfer function. Such a function is thought to be computable on either biological or mechanical hardware, thereby rendering the internal workings of thought irrelevant. This technocratic impulse has been used to sell AI & ML products as “magical” solutions, capable of ushering in Utopian futures. This viewpoint began with the foundation of computer science itself, as metaphors for computational processes were adopted without adequate grounding in the philosophy of mind. Computers were given attributes of human cognition as a teleological basis for investment in these technologies. Our current situation sees the world’s accumulated media scraped for so-called “knowledge bases,” in many cases reinforcing cultural biases, ignoring creator rights, and consuming energy resources... all to create pastiches of existing art. As critical consumers, we should evaluate each novel technology for the social, cultural, and political assumptions that underlie their functioning. This paper will take steps towards that goal by analyzing Google’s MusicLM as a test case. The musical artists Delia Beatriz, Nao Tokui (interviewed elsewhere in this issue) and Moisés Horta Valenzuela will be used as exemplars of creative engagement with ML.

### **Metaphor considered harmful**

The field of Artificial Intelligence sprang to life at a 1956 workshop organized at Dartmouth by mathematician John McCarthy<sup>i</sup>. McCarthy coined the term to differentiate himself from the cybernetics of Norbert Wiener<sup>ii</sup>, an individual with whom he had profound personal, professional, and political disagreements<sup>iii</sup>. At its inception, Wiener’s cybernetics studied communication and control through goal-directed, teleological behaviors — self-regulation through feedback and feed-forward loops — across electronic, mechanical, and biological systems. By contrast, McCarthy’s AI was focused on the science and engineering of “intelligent machines.” This central metaphor has become problematic in our current era, where terms like “intelligence” and “learning” are deployed to make these technologies seem more capable than they actually are<sup>iv</sup>.

As an example of how these metaphors function, consider how the computer desktop GUI acts as a visual metaphor for computational processes<sup>v</sup>. The use of a physical bin to collect recyclables is mapped onto the desktop-level task of deleting a file, itself a placeholder for the underlying computations necessary to accomplish this task at the level of disk storage. In this way, metaphors allow us to describe hidden processes in terms of the familiar and accessible<sup>vi</sup>. In doing so, they mediate our understanding, imposing a set of conceptual and linguistic relations. Computer metaphors encourage us to think in terms of learning, knowing, understanding, thinking,... even consciousness and sentience.

Such metaphors go back to the very genesis of computer science. In 1943 McCulloch and Pitts described a system of nerves as a “nervous net,” which they defined using properties and processes then known to physiology. They proposed that this net (i.e. the human brain) was congruent with a Turing machine<sup>vii</sup>. In a remarkable conclusion, the authors then posit (quite without justification) “ultimate psychic units” they termed “psychons”<sup>viii</sup>, hence reducing the processes of psychology to the same binary logic they had applied to physiology. McCulloch and Pitts described diseases of the mind as deviations from perfect operation. The task of psychiatry was simplified to recognizing, categorizing, and correcting such logical errors. There’s no room for neuro-diversity in such a rationalization.

John von Neumann’s pivotal “First Draft of a Report on the EDVAC” also described components of the computer as “neurons in the human nervous system” while larger components (input and output blocks, for example) were labeled “organs,” a term that has not persisted in the lexicon<sup>ix</sup>. Computer storage was metaphorically described as “memory,” without any further justification. Is fixing a bit in a circuit equivalent to the malleable constructs networked in nerve clusters? Is the operation of data retrieval the same as recalling a memory from the associative field of emotion, signification, and imagination? A moment’s consideration should challenge such notions. But the metaphors were compelling to those relatively unfamiliar with philosophy (or Proust for that matter). Soon the terms “thinking machine” and “computer brain” were fixed in book titles and course notes. The reductionist idea that symbolic thought could be modeled by the operations of switches became commonplace. If indeed the mind was this simple, then computational wonders would no doubt be accomplished in due course. Utopian proposals for AI made possible enormous budgets for “the biggest and most expensive computers available” and thus were successful at the level of building infrastructure<sup>x</sup>. But they also propagated the mistaken idea that a computer might possess consciousness, without due consideration of how such epiphenomena might arise.

A metaphor can be a useful shortcut, a hypothetical construct designed to aid understanding. But when we forget that we are using a metaphor, when we mistake an abstraction for a concrete thing in the world, we commit the fallacy of reification<sup>xi</sup>. Rather than helping us understand computer systems, this error hampers our capacity to engage critically with technology, while helping corporations sell technological products to us as magic<sup>xii xiii</sup>. Arthur C. Clarke's third law, so beloved of technocrats, posits that "any sufficiently advanced technology is indistinguishable from magic"<sup>xiv</sup>. This idea is deployed to venerate the devices, services, and personas that drive the technical industries, while conjuring clouds of obfuscation around their production and operation. A more critical reading of Clarke's third law might instead suggest that anything presenting itself as magic is, in fact, a technological system. Our task then is to interrogate that system's intended purpose and benefactors; its social, cultural, and political functions; the method by which it was produced; the contextual framework in which it was generated; and so on.

Belief in computer thought and sentience has certainly captured the imagination, producing science-fiction scenarios such as those represented in *Alphaville* (dir. Jean-Luc Godard, 1965), *2001: A Space Odyssey* (dir. Stanley Kubrick, 1968), *Colossus: The Forbin Project* (dir. Joseph Sargent, 1970), and *The Terminator* (dir. James Cameron, 1984). From these fictions we are led to believe that the intelligence of machines might cause harm to mankind. Though this at first appears contrary to the Utopian precepts of AI, both formulations (Utopian and Dystopian) share an error in common: they ascribe intentionality to the artificial intelligences themselves. Our fictions have trained us to think that malevolent machines will cause mischief, when in truth it is the person/corporation behind the machine that we must question. We should not fear Skynet but instead guard against Cyberdyne Systems.

### **Ethical concerns**

To illustrate the complex ethical quandaries around AI and ML, we'll consider the domain of music generation, Google's MusicLM product in particular<sup>xv</sup>. MusicLM is a machine learning system for generating music tailored to user specifications. You need merely type a phrase describing what you wish to hear and the engine will generate a result. The complex architecture of MusicLM consists of multiple components, each incorporating voluminous data. MuLan, the largest of these subsystems, has been trained on 44 million 30-second sound clips, each mapped to corresponding text annotations, which were scraped from metadata, comments, and playlist data<sup>xvi</sup>. Where exactly these clips originate is not explained in the MuLan paper. But the sample size, along with the description "internet music videos," strongly implies that the data was taken

from YouTube. It would be impossible to quantify the immense human labor that went into producing the 44 million original works posted freely to that video host. The creators of this “data” received neither credit nor compensation, were not asked for their consent, and wouldn’t even realize that their original work was being used in MuLan (unless they read the paper in question). This (mis)use of information opens up ethical questions around permissions, rights ownership, and copyright theft that deserve considerable attention. But it is symptomatic of the technocratic ethos that these topics are not addressed in such papers. AI ends are pursued without considering the means.

While MuLan’s text annotations are acquired through scraping practices, this is not always the case. Indeed, human labeling for AI datasets is a big business. Google, Meta, and Microsoft have outsourced labeling duties to firms such as Sama, based in San Francisco. A recent report in *Time* highlighted how Sama offshore employees earned between \$1.32 and \$2 per hour while working to label data for OpenAI’s ChatGPT<sup>xvii</sup>. This low wage demonstrates how ML algorithms are predicated on labor exploitation.

The training of systems like MusicLM and ChatGPT also have high environmental cost. The initial training of GPT-3, the model underlying the initial ChatGPT release, consumed 1.3 GWh of electricity, enough to power 120 homes for a year, and in the process released an estimated 552 tons of CO<sub>2</sub><sup>xviii</sup>. The training of GPT-4 used an estimated 7.5 GWh, enough to power 700 homes for a year. Google’s AI projects consume enough electricity in one year to power every home in the city of Atlanta for the same period<sup>xix</sup>. These numbers will grow by an estimated 8 GWh per year, as models increase in size and scope.

We should also be wary of how data is gathered for these training models, from where data is sourced and what constraints are (or are not) applied. It has been found that white supremacist, misogynistic, and ageist views are over-represented in training data scraped from the internet, and these biases tend to get amplified even further in ML data sets<sup>xx</sup> <sup>xxi</sup>. (Recall that ML systems don’t understand meaning, and hence cannot make any sort of judgment about the data they are processing.) As this material is repurposed for curious users, the risk of real-world harm (physical, economic, social and emotional) for marginalized communities increases.

Generative Adversarial Networks (GAN) also tend to replicate the biases of their programmers and the knowledge bases that they are fed. For example, when Snapchat beautifies a woman’s face, it leaves light-skinned subjects alone but lightens darker skins<sup>xxii</sup>. When Stable Diffusion is asked to picture an “ambitious CEO” it includes only men, but given the phrase “supportive CEO” it allows women into the solution set<sup>xxiii</sup>. These biases extend to musical applications where gender imbalances have been detected in models developed for music

recommendation systems. Problems of bias and limited diversity have been flagged at both the level of the dataset and algorithm design in music information retrieval<sup>xxiv</sup>.

### **From pastiche to experimentation**

We've demonstrated that ML systems exploit labor, consume extensive resources, and readily reinforce societal biases of gender, race, and other attributes. But what of the resulting music? To a trained ear, and perhaps even to the casual listener, the results are bland and imitative. We might best describe the results as *pastiche*<sup>xxv</sup>, for two reasons. First, because the outputs are always backward-looking, predicated on recombining previous works. The inference engines of ML embed processes of replication, interpolation, and extrapolation. But the outputs reflect patterns present in the original training data, even in the simple terms of musical attributes (volume, timbre, tempo, etc.). Often results will tend towards the mean value profiles for each of these that are represented across the training data.

Second, in stark contrast to human-created music, ML models generate music from material that's divorced from sociopolitical context. If all we'd ever had was ML, we'd never have heard a jazz or blues reaction to the oppression of an African-American underclass (hence Billie Holiday would never have sung "Strange Fruit"). There would never have been a punk rock excoriating the English Queen's Jubilee Year (hence the Sex Pistols would never have belted out "God Save The Queen"). A musical knowledge base contains no extra-musical information that might generate novel reactions to quotidian reality.

So, yes, the normative result of systems like MusicLM are bland pastiche. But do we have alternatives? Can experimental sound practices transcend this *modus operandi*? Can ML techniques be subverted to achieve something more interesting than "melodic techno" and "relaxing jazz" (to use two examples from the demonstration website)? To begin answering these questions, we should first define "experimental" music by invoking Cage's declaration that "an experimental action is one the outcome of which is not foreseen"<sup>xxvi</sup>. The most extreme method to accomplish this would be completely stochastic, but George E. Lewis points out that a focus on unique and spontaneous chance operations in every moment eliminates personality, narrative, memory, and history from the experience<sup>xxvii</sup>. Such is not our goal. With Lewis, we hold music to be an expression of the lived experience of the players involved; collaborative improvisation is an assertion of their agency. It's on these principles that Lewis created Voyager, an interactive computer music environment that operates as a virtual improvising orchestra<sup>xxviii</sup>. Voyager was developed by combining principles from 1980s AI and 1950s cybernetics with socio-musical networks of free improvisation<sup>xxix xxx</sup> more recently expanded with AI/ML

techniques<sup>xxxii</sup>. Here the machine and human are equal collaborators within a context that involves unexpected improvisatory extrapolations on the basis of personality, narrative, memory, and history. Here, ML techniques provide a rich starting point for experimental sonic practices<sup>xxxiii</sup>, reducing the likelihood of pastiche.

Subsequently a growing contingent of artists and researchers are developing, appropriating, and subverting ML technologies to generate exciting musical possibilities. Furthermore, they are doing so in a manner that addresses the ethical shortcomings of these technologies.

### **A humanistic turn in AI/ML-driven art**

In the work of Tokyo-based artist, DJ, and researcher Nao Tokui we see an approach to the application of AI that is concerned with expanding human creativity. In his interview for *Resonance*, available in this issue, Tokui describes an artistic practice that exists on the borderline between chaos and control where the AI tools that he creates, and then applies in his own work, allow new artistic possibilities to come into being. In 2015 Tokui worked with Brian Eno on the “Generative Film” for his album *The Ship*. Reflecting on his work with Eno, Tokui recalls:

When I first worked with Brian Eno, he said, “I’m not really interested in AI.” Rather, he was interested in “Artificial Stupidity.” If what we’ve been working on up to this point is correct, then there are new discoveries and ideas among the deviations, mistakes, and foolishness. I think being able to accept that is creativity.<sup>xxxiii</sup>

For Tokui, and Eno, the metaphor of Artificial Intelligence is not just lacking but boring. It’s too perfect, too predictable. Interest and creativity are to be found in mistakes, deviations, and outliers. For Tokui, it is the playful subversion of AI that extends the possibility of human creativity. New creative potentialities are opened up through the kinds of happy accidents that can only result from the deliberate subversion, or what he terms “misuse”, of AI tools. This principle of playful subversion is embodied in the works he discusses in our interview and is particularly evident in his live audiovisual improvisations. Tokui invites us to stop thinking about AI as an intelligent agent and instead think of it as a tool that can be subverted to uncover otherwise impossible creative possibilities.

On her 2022 record *The Long Count*, the artist known as Debit (Mexican-American producer Delia Beatriz) used ML techniques to create novel digital instruments. Her source material was

recordings of precolonial Mayan wind instrumentation from the archive of the Mayan Studies Institute at the Universidad Nacional Autonoma de Mexico<sup>xxxiv</sup>. She spent a month standardizing the data set for CREPE, “a deep convolutional neural network which operates directly on the time-domain audio signal to produce a pitch estimate”<sup>xxxv</sup>. The resulting frequency data for each sample allowed virtual instruments to be built as composites of the archival material<sup>xxxvi</sup>. These instruments formed the basis of compositions that function as a kind of electroacoustic archaeology. Because we don’t know what Mayan music sounded like, as no written music survived the Spanish conquest, the resulting pieces are both speculative and anachronistic in nature, neither part of the past nor the present but haunted by indistinct, shared memories.

The instruments I sampled could have once been used to communicate with Jaguars or jungle cats, which were influential in their culture. Mayans did want to use these instruments for interspecies communication, as well as possibly time-hacking, ceremonies or communing with the divine. They were used for more than just music, but we don't really know to what effect those instruments were employed<sup>xxxvii</sup>.

There is a process of decolonization at play here, but also a kind of speculative restoration in the reclamation of Mayan instrumental sounds in a modern electroacoustic context. The result is a music that is both new and ancient, subverting both the anachronistic processing inherent in ML and the predictable nature of the music generally so produced. The metaphor of AI as something magical is undermined by the disturbing nature of the resulting drones and howls. But there is some new “magic” in these results. Not the magic of a promised technocratic Utopia, but the magic of an evocative dialogue between speculative futures and erased pasts.

A second Mexican-American artist who has transformed ancient sounds into novel configurations is Moisés Horta Valenzuela. Based in Berlin and recording as Hexorcismos, his 2020 album *Transfiguración* used Antonio Zepeda’s 1982 album *Templo Mayor* as source material. Zepeda was one of the first to record with ancient Mesoamerican and ethnic Mexican instrumentation. This album is important to Valenzuela, who claims it as a site of memory and resistance<sup>xxxviii</sup>. Transforming this pre-Hispanic material using a GAN allows the artist to address his own “migratory modes of sound production.” The result is a combination of ancient instruments, a classic 1980s recording of personal significance to the artist, and contemporary electronic sound production. Valenzuela pursues this course with vigor.

I feel these works, the album and this mix, is about mnemonic resistance, about memory. The colonial processes that happened in Mexico and in most post-colonial

countries is that of epistemicide (erasure of indigenous cosmovisions). The remnants of the cultures that have survived throughout the years in the face of capitalistic modernity have adapted, either through embracing or hybridizing with the colonial cultures<sup>xxxix</sup>.

Valenzuela enacts a critical engagement with the conceptual and technological apparatus of AI as developed within the military-industrial complex of the Global North. This allows him to reclaim an erased musical past, albeit in a different manner from Beatriz. Whereas Beatriz used neural net techniques to produce digital instruments with which she could compose, for Valenzuela the GAN itself becomes the tool of transformation. Indeed, he goes so far as to claim the GAN is “a technoshamanistic conceptual framework,” hence reclaiming the magic inherent in the system as “a kind of alchemist tool.” Here again, as with both Beatriz and Tokui, we find that the subversion of AI opens rich artistic vistas that would otherwise remain closed.

## **Conclusion**

The Utopian claims of AI supporters, from the earliest computer scientists to today’s technocrats, are predicated on functionalism, the doctrine that computer “thought” is the equivalent of human thought, so long as each produces the same responses to input. The internal functioning of these “minds” is irrelevant, as explicitly claimed by Turing when proposing his famous Imitation Game<sup>xl</sup>. But this classical computational theory of mind (CCTM) is predicated on the fallacy of reification, taking as concrete what should remain as metaphor. CCTM can only maintain self-consistency by denying several important qualities of human thought. Primary among these is the fact that an algorithmic process within a Turing machine is not “about” anything. By way of contrast, human thought is always intentional, thought *about* something, responding to and generating *meaning*. A computer has no knowledge of meaning, no matter how large the data set<sup>xli</sup>. Indeed, it would be entirely possible for a machine to pass the Turing test without needing to “think;” it need only imitate sufficiently well (as the name of the test admits).

Taking a critical stance against functionalism is important for reasons already elucidated. Unless we examine the claims of AI and ML with a watchful eye, we might assume that such tools have only positive value, designed to improve our lives. We’d remain unaware of the sociopolitical biases inherent in such systems, the capitalist excesses of resource consumption and labor exploitation inherent in their very materiality.

If a computer’s task is to answer questions, the artist’s role is to question answers. As exemplars of such critical thinking, we’ve presented Tokui, Beatriz, and Valenzuela. All three are



pushing the bounds of AI art. Tokui's work expands our understanding of what human creativity can be, through intentional subversion and misuse of AI. Beatriz and Valenzuela's work subverts these same technologies to reclaim Mexican (pre-) history from excision. None of these artists aim to preserve in amber any rose-colored past, but rather to manifest new aesthetic expressions from rich material otherwise overlooked. They use neural networks not to supplant or better human thought, but to create spectacular imaginings that subvert the "machine intelligence" metaphor. Beatriz and Valenzuela do so by rooting their processes in a deep respect for musical cultures that have been damaged and erased while Tokui discards the "machine intelligence" metaphor and foregrounds instead the expansion human creativity in all its messy glory. Their work manifests an ethos of opposition, challenging those who would pillage the creative work of a million artists for the sake of pastiche.

## Notes

---

<sup>i</sup> John McCarthy. *Defending AI research: a collection of essays and reviews*. (CSLI Publications, Stanford, 1996).

<sup>ii</sup> McCarthy, *Defending AI*, 73.

<sup>iii</sup> McCarthy, *Defending AI*, 45-46.

<sup>iv</sup> Kate Crawford, *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (Yale University Press, 2021).

<sup>v</sup> David Benyon and Manuel Imaz, "Metaphors and models: conceptual foundations of representations in interactive systems development," *Human-Computer Interaction* 14.1-2 (1999): 159-189.

<sup>vi</sup> George Lakoff and Mark Johnson, *Metaphors We Live By* (University of Chicago Press, 1980).

<sup>vii</sup> Warren S. McCulloch and Walter Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics* 5 (1943), 129.

<sup>viii</sup> McCulloch and Pitts, 132.

<sup>ix</sup> John von Neuman. "First Draft of a Report on the EDVAC," *IEEE Annals of the History of Computing* 15.4 (1993, original 1945), 27-43.

<sup>x</sup> Thomas Haigh, "Conjoined twins: Artificial intelligence and the invention of computer science," *Communications of the ACM* 66.6 (2023), 33-37.

<sup>xi</sup> Whitehead called this the "Fallacy of Misplaced Concreteness" in Alfred North Whitehead, *Science and the Modern World* (New York: Pelican Mentor / New American Library), 52.

<sup>xii</sup> Matthew Gwynfryn Thomas and Djuke Veldhuis, "Is artificial intelligence magic?" *Sapiens*, 12 December 2019.

<sup>xiii</sup> John Herrman, "The AI magic show," *New York Magazine*, 8 January 2023.

<sup>xiv</sup> Arthur C. Clarke, "Hazards of Prophecy: The Failure of Imagination," *Profiles of the Future: An Inquiry Into the Limits of the Possible*, revised edition (New York: Harper and Row, 1973), 21.

<sup>xv</sup> Andrea Agostinelli, Timo I. Denk, Zalán Borsos, Jesse Engel, Mauro Verzetti, Antoine Caillon, Qingqing Huang, Aren Jansen, Adam Roberts, Marco Tagliasacchi, Matt Sharifi, Neil Zeghidour, and Christian Frank, "MusicLM: Generating music from text," arXiv:2301.11325, 2023.

<sup>xvi</sup> Qingqing Huang, Aren Jansen, Joonseok Lee, Ravi Ganti, Judith Yue Li, and Daniel P. W. Ellis, "MuLan: A joint embedding of music audio and natural language," arXiv:2208.12415, 2022.

<sup>xvii</sup> Billy Perrigo, "Exclusive: OpenAI used Kenyan workers on less than \$2 per hour to make ChatGPT less toxic," *Time* (website), 18 January 2023.

<sup>xviii</sup> Molly Loe, "Heating up: How much energy does AI use?," *TechHQ*, 24 March 2023.

<sup>xix</sup> Josh Saul and Dina Bass, "Artificial intelligence is booming: So is its carbon footprint," *Bloomberg*, 9 March 2023.

<sup>xx</sup> Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell, "On the dangers of stochastic parrots: Can language models be too big?," *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623. 2021.

<sup>xxi</sup> Timnit Gebru, "Race and gender," *The Oxford Handbook of Ethics of AI* (Oxford University Press, 2020): 251-269.

<sup>xxii</sup> Niharika Jain, Alberto Olmo, Sailik Sengupta, Lydia Manikonda, and Subbarao Kambhampati. "Imperfect ImaGANation: Implications of GANs exacerbating biases on facial data augmentation and snapchat face lenses," *Artificial Intelligence* 304 (2022), 103652, <https://doi.org/10.1016/j.artint.2021.103652>.

<sup>xxiii</sup> Vongani H. Maluleke, Neerja Thakkar, Tim Brooks, Ethan Weber, Trevor Darrell, Alexei A. Efros, Angjoo Kanazawa, and Devin Guillory. "Studying bias in GANs through the lens of race" in S. Avidan, G. Brostow, M. Cissé, G.M. Farinella, T. Hassner, eds., *Computer Vision – ECCV 2022*. Lecture Notes in Computer Science 13673 (Switzerland: Springer, 2022). [https://doi-org.proxy.lib.ul.ie/10.1007/978-3-031-19778-9\\_20](https://doi-org.proxy.lib.ul.ie/10.1007/978-3-031-19778-9_20).

<sup>xxiv</sup> Sturm, Bob LT, Maria Iglesias, Oded Ben-Tal, Marius Miron, and Emilia Gómez, "Artificial intelligence and music: open questions of copyright law and engineering praxis," *Arts* 8.3 (2019): 115.

<sup>xxv</sup> In using the term "pastiche" we generously allow that the processes are designed to honor the original musical intentions. Otherwise, the results would more properly be described as parody.

<sup>xxvi</sup> Cage, John. *Silence: Lectures and Writings* (Wesleyan University Press, 2012).

<sup>xxvii</sup> Lewis, George E. "Improvised music after 1950: Afrological and Eurological perspectives," *Black Music Research Journal* (1996): 91-122.

<sup>xxviii</sup> Lewis, George E. "Too many notes: Computers, complexity and culture in 'Voyager'." *Leonardo Music Journal* (2000): 33-39.

<sup>xxix</sup> Lewis, George E. *A Power Stronger Than Itself: The AACM and American Experimental Music* (University of Chicago Press, 2019).

<sup>xxx</sup> Lewis, George E. "Is our machines learning yet?," *Machinic Assemblages of Desire: Deleuze and Artistic Research* 3 (2021): 115.

<sup>xxxi</sup> "George Lewis Residency | RNCM Musicians and an Enhanced Voyager," Royal Northwestern College of Music, last modified June 1, 2023, <https://www.rncm.ac.uk/research/research-centres-rncm/prism/prism-news-and-events/george-lewis-residency-rncm-musicians-and-an-enhanced-voyager/>.

<sup>xxxii</sup> Stephen Roddy, "Creative machine-human collaboration: Toward a cybernetic approach to artificial intelligence and machine learning techniques in the creative arts," *AI and the Future of Creative Work*, ed. Michael Filimowicz (UK: Routledge Press, 2023), 177-78.

<sup>xxxiii</sup> "Surfing in the middle ground between control and chaos is the key to dealing with AI," *Wired Japan* (website), last modified June 15, 2019, <https://wired.jp/series/away-from-animals-and-machines/chapter7-1/>.

<sup>xxxiv</sup> "The Long Count," *Pitchfork Music* (website), last modified February 23, 2022, <https://pitchfork.com/reviews/albums/debit-the-long-count/>

<sup>xxxv</sup> Jong Wook Kim, Justin Salamon, Peter Li, Juan Pablo Bello, "CREPE: A convolutional representation for pitch estimation," 17 Feb 2018, arXiv:1802.06182v1.

<sup>xxxvi</sup> Martyn Pepperell, "Cyclic revolution: Debit's radical new album reconstructs Mayan musical history," *Mixmag* (website), 4 March 2022, <https://mixmag.net/feature/debit-interview-the-long-count-mayan-music>.

<sup>xxxvii</sup> Zach Sokol, "Meet DEBIT, experimental producer and Azealia Banks co-conspirator," *i-D* (website), 30 March 2022, <https://i-d.vice.com/en/article/bvnaa3/debit-the-long-count-interview>

<sup>xxxviii</sup> Moisés Horta Valenzuela, "Hexorcismos – Transfiguración" (website), <https://moiseshorta.audio/obra/transfiguracion/>.

<sup>xxxix</sup> Peter Kirn. "E Transfiguración: decolonizing AI, in Hexorcismos' shamanistic music and art", *Create Digital Music* (website), July 2020, <https://cdm.link/2020/07/transfiguracion-decolonizing-ai-in-hexorcismos-shamanistic-music-and-art/>.

<sup>xl</sup> A.M. Turing, "Computing machinery and intelligence," *Mind* 49 (1950): 433-460.

<sup>xli</sup> There are many other objections to functionalism, but not all need to be rehearsed here. But consider that a computer does not have emotions nor bodily sensations, hence cannot enter experiential states in which perception of one's own state is vital. Belief in a computational theory of mind often infers

Roddy, S., & Parmar, R. (2023). AI/ML in the Sonic Arts - Pitfalls and Pathways. *Resonance: The Journal of Sound and Culture*, 4(4). <https://doi.org/10.1525/res.2023.4.4.399>

---

the possibility that brain and mind can be split, a proposal that has no basis in fact (even if it makes for intriguing science-fiction scenarios).