2 Creative Machine-Human Collaboration

Toward a Cybernetic Approach to Artificial Intelligence and Machine Learning Techniques in the Creative Arts

Stephen Roddy

Recent Developments in AI Applied to Art

There have been giant leaps forward in the application of cutting-edge artificial intelligence (AI) and machine learning (ML) and more specifically deep learning (DL), technologies in creative and cultural industries in recent years. These changes have been felt in the visual arts and music in particular. After a brief technical overview later, the impact of these technologies on the contemporary landscape in both fields is explored with reference to their historical antecedents.

Audry's succinct description describes ML algorithms as processes that underpin "computational systems that are biologically inspired, statistically driven, agent-based networked entities that program themselves" (Audry, 2021). Within this somewhat broad definition, we find a myriad of ML model architectures that can be adapted to a wide range of tasks in the creative arts. Large language models with transformer-based architectures like T5 (Raffel et al., 2020) and the GPT series (Brown et al., 2020), have revolutionized the production of human-like text and AI-driven chatbots, ChatGPT and Bard in particular, are whipping up a flurry of fevered speculation about the nature of machine intelligence (Warzel, 2023). Generative models for visual production like Midjourney (Midjourney, 2023), DALL-E 2 (OpenAI, 2021a), Imagen (Google, 2022a) and Parti (Google, 2022b) generate a wide range of imagery from simple text prompts. DALL-E 2 and Imagen employ diffusion models, which operate by destroying their training data through the addition of Gaussian noise before learning to recover the original data through a process of reversal (Dhariwal & Nichol, 2021). Stable Diffusion v1 (Rombach et al., 2022) also shows promising results for image generation and modification. Image modification converts simplistic sketches into detailed artworks, something that Nvidia's GauGAN2 can achieve in close to real time (Park et al., 2019). Both GauGAN2 and Parti employ the Generative Adversarial Network (GAN) architecture introduced by Goodfellow et al. (2014). Other comparable

DOI: 10.4324/9781003299875-2

models include StyleGAN3 (Karras et al., 2021) and VQGAN (Esser et al., 2021), a variant of which is used in Parti. These models can be combined with CLIP (OpenAI, 2021b) to handle text prompts for image generation tasks.

There have been developments in the production of sound and music too. MuseNet adopts a similar approach to GPT-2 and is trained on and designed to produce MIDI files that can be synthesized or mapped to a sonic output later (Payne, 2021). The Magenta project has developed Variational Autoencoders or VAE (Roberts et al., 2018) for generating music that is designed to work with MIDI data. They have also developed a technique branded differentiable digital signal processing or DDSP for transforming the sounds of musical instruments (Engel et al., 2020). This approach works on raw audio waveforms, as opposed to MIDI files, in a similar fashion to OpenAI's WaveNet (Oord et al., 2016), which was a breakthrough model for raw audio generation in 2016. A similar approach is employed by OpenAI's Jukebox (Dhariwal et al., 2020), which uses a VAE architecture trained on raw audio files to generate music across a wide variety of genres. There is now an ever-expanding range of websites and applications that use ML technologies to creative ends.

Artists have been working with AI since at least the late 1960s when Harold Cohen began work on AARon, a rule-based system from image generation that simulated cognitive primitives thought by Cohen to underpin drawing and painting (Cohen, 1982). Lejaren Hiller and Leonard Issacson's Illliac Suite, a 1957 string quartet in four movements, is thought to be the first computer-assisted composition (Hiller & Isaacson, 1957). The pair devised increasingly complex compositional methods for each movement beginning with some simple harmonic rules for the first movement and ending with the application of Markov Chains for the fourth. In a similar fashion to Cohen, David Cope began work on his experiments in musical intelligence (EMI) system in 1981 (Cope, 1992), which generated music by analyzing musical input data for core signifiers of the works that could be retained during the recombination of musical sequences to create novel works.

The recent ML boom, driven by the rise in availability of computing hardware like high-performing GPUs and cloud computing infrastructures, the availability of rich big datasets, and the improved effectiveness of DL techniques, has led to a flood of new artists working with AI and ML systems. The techniques adopted during this boom period stand in stark contrast to earlier artists like Harold Cohen and David Cope who worked with rule-based AI, sometimes referred to as Good Old Fashioned AI or GOFAI or artists working in a visual medium like Frieder Nake and Herbert Franke who were interested in computational systems more generally.

Ahmed Elgammal (2019) argues that the work being produced in this epoch is a form of conceptual art because the focus is on the creative process rather than the piece produced and the collaboration between artist and machine required to produce it. Miller (2019) describes how Alexander Mordvintsev's Deep Dream marked a turning point for AI art. This system creates dream-like imagery by applying a convolutional neural network to amplify and iterate upon patterns in input images. Important works in this epoch include Mike Tyka's (2018) Portraits of Imaginary People, which used a GAN trained on photos of faces collected from Flickr.com to generate new faces of people who do not exist. Mario Klingemann's (2017) Imposture Series meanwhile produced works with a GAN trained on varied combinations of stick figures and imagery harvested from the internet. Refik Anadol's (2016) Machine Hallucination series uses huge sets of visual data (e.g., photos from the International Space Station or photos of natural phenomena and landscapes) to train GAN-type models and then renders his pieces as walks through the model's latent space: the compressed representation of the features from the original data that has been learned by the model. Memo Akten's (2017) Hello World trains a VAE on a live video feed allowing the user to change model hyperparameters in real time as it trains. Pindar Van Arman's Cloud Painter project uses DL techniques to teach robots to paint like people (Miller, 2019). These changes have impacted markets in the creative and cultural industries. In 2018, the arts collective Obvious sold a visual artwork created with a GAN they sourced on GitHub. It sold for $432,000 at Christie's, which was particularly problematic given that the GAN in question was written by AI artist Robbie Barat, a teenager at the time, who was excluded from the profits of the sale (Miller, 2019; Flynn, 2018).

A notable early application of neural networks in the creative arts is Rebecca Feibrink's Wekinator (Fiebrink et al., 2009), which uses neural networks to learn bespoke mapping strategies from real-time control data inputs to the parameters of a given multimedia system like a synthesizer. In 2017 Taryn Southern and Benoît Carré's SKYGGE project both produced albums that incorporated ML techniques. Southern's I AM AI made use of IBM's Watson Beat and tools provided by amper.com (Jancer, 2018), while SKYGGE's Hello World worked with tools created by flow machines to produce his pieces (SKYGGE, 2017). That same year Dadabots began releasing works that involved WaveNet-style raw audio models trained on a variety of music but particularly Black Metal (Zukowski & Carr, 2018). In 2018 Ash Koosha explored AI on Return 0 (Cardew, 2018), and in 2019 Auxuman (Fry, 2021), an AI collective co-founded by Koosha, began monthly musical releases. Lee Gamble (Quarshie, 2021), Mouse on Mars (Sherburne, 2021) and Arca (Darville, 2020) have incorporated a variety of ML techniques into their music and Hexorcismos' Transfiguración devised a novel approach to synthesizing audio with GANs (Kirn, 2020). Bob Sturm and Oded Ben-Tal are working with real-world traditional music practitioners to produce novel new musical scores using ML/DL techniques (Miller, 2019). Another highly original approach is Holly Herndon's Holly+, an AI-driven vocal deepfake tool that allows users to make music their own music using her voice (Holly, 2021).

While the works of musical artists engaging with ML technologies can be contextualized within a larger tradition of computer music and computational art more broadly (Miller, 2019), neural networks and ML techniques only entered the picture in a music composition context in the late 1980s with the publication of pioneering work by Lewis (1988) and Todd (1988). Following this, David Tudor's (1995) Neural Synthesis N° 6–9, was a collection of pieces composed with a hardware synthesizer designed for Tudor by Forrest Warthman, Mark Holler and Mark Thorson. The synthesizer was built around an analog neural network microchip: the Intel 80170NX neural processor or electronically trainable analog neural network (ETANN) (Kuivila, 2004). Eck and Schmidhuber (2002) would be the first to apply long short-term memory models (LSTMs) in a musical context and with the DL boom of the early 2010s a wide range of ML applications in both music and the digital arts emerged. While the application of ML to the arts is a more recent development, its focus and development have nonetheless been shaped and constrained by historical forces, which have defined how we conceptualize, design and interact with intelligent machines.

Problematic Foundations

The human relationship to machines and the concept of intelligent machines, in particular, has long been shaped by a central anxiety about the eventual replacement of humans by machines. In *Machines Who Think: A Personal Inquiry Into the History and Prospects of Artificial Intelligence*, Pamela McCorduck lays out a scathing criticism of, among other things, the master-servant dynamic that underpins a lot of our thinking about intelligent machines (McCorduck, 1979). She highlights how even the word "robot" is coded with these values having first appeared to describe the factory-fabricated android servants in Karel Čapek's play, *R.U.R*. (Rossumovi Univerzální Roboti—Rossum's Universal Robots). The term was dreamed up by Čapek's brother Josef and is derived from the Czech word for "servitude." It is also worth noting that the titular robots of the play revolt against their human masters, bringing about the extinction of humanity. The anxiety that one day intelligent machines we have created and indentured will rise up and turn the tables on their human masters has been a theme in the cultural and literary depictions of the robot since its very inception. Even earlier in 1872 Samuel Butler's *Erewhon* (see Butler, 2015), specifically the three chapters that constitute *The Book of the Machine*, introduced the idea that machines might become conscious through a Darwinian process of natural selection and thus learn to self-replicate also. While he would later come to believe that a mechanistic model of the organism would render it incapable of consciousness (Breuer, 1975), his writings in *The Book of the Machine* suggested that intelligent machines were to be feared as a dangerous threat, destined to gradually supersede and replace humanity as a dominant force.

The AI takeover would become a recurring theme in 20th-century science fiction as writers and creators who, following Butler, explored and expanded upon the idea across a wide range of media. While Isaac Asimov famously rejected this notion, the theme appears in highly influential works by Stanisław Lem (Lem, 2021), Philip K. Dick (Dick, 1968) and Harlan Ellison (Ellison, 1967) that would in turn influence landmark works by filmmakers like Stanley Kubrick, Ridley Scott and James Cameron (Roberts, 2016).

The concept of the "technological singularity" was introduced by John von Neumann (Shanahan, 2015) and expanded by Alvin Toffler (Toffler, 1970). It developed into a cohesive vision in the writings of Vernor Vinge (Vinge, 1993) and Ray Kurzweil (Kurzweil, 2005) and would eventually be given a cursory mathematical formalization by Nick Bostrom (2014). The idea generally goes that our continually improving computing technologies will soon lead to the emergence of runaway superintelligences that radically reshape society and humanity. The idea is not without its critics. Pein (2018) provides criticism of both Vinge and Kurzweil, pointing out the lack of a sound scientific basis for their predictions about the singularity and attacking the cult-like nature of the movement that has grown up around those predictions. Benthall (2017) uses a Bayesian model to demonstrate that the probability of Bostrom's intelligence explosion actually happening, based as it is on the advancement of recursively self-improving AI algorithms, is negligible given the importance of hardware and data to the growth of intelligence. Kurzweil's singularity is one wherein, rather than replacing humans, superintelligent AI will enhance and improve humanity as it integrates more closely with it. It eventually leads to humans uploading scanned copies of their brains (and therefore minds according to Kurzweil) into these superintelligent machines thus achieving a form of techno-mediated immortality. By contrast, for Vinge, superintelligent machines are a dangerous threat and, in the best-case scenario, the one in which they don't simply wipe out all human life, they must be bound as "godlike slaves" to the will of their human masters. This is echoed by Bostrom who sees it as imperative that we solve the "control problem" to prevent an existential catastrophe. His various methods for controlling intelligent machines include containment, stunting, self-destruct scenarios and a kind of machine eugenics that might select for domesticity and normative alignment. Vinge and Bostrom's singularity scenarios are underpinned by a sense of anxiety that humans will be replaced by machines and as such humanity must double down in its role as subjugator, ruthlessly enforcing the master-servant dynamic that already defines so much of the relationship of humans to machines.

As McCorduck points out, this dynamic has been with human thought for a long time. We see early descriptions of automata, mechanical humans, in the 4th-century BCE Daoist text the *Liezi* attributed to the 5th-century philosopher Lie Yukou (Richey, 2011). The artificer Yan Shi creates a wooden automaton that can dance and sing. The automaton attracts the ire of King Mu when it breaks from the established protocols of the court and begins to wink and advance toward the ladies of the court causing Yan Shi to take him apart. Both a master-servant dynamic and a level of anxiety with regard to the robot breaking from its assigned role in this hierarchy are on display here. It is also noteworthy that this early depiction of an automaton presents it as a creative being: a skilled singer and dancer. This is of course because it was created to entertain the powerful, in this case, King Mu, to win favor or reward for his creator. In the Iliad, we again see the master-servant dynamic at play when Homer describes Hephaestus' servants as intelligent and articulate young women, wrought from gold and portrayed as supportive of their master's every move (Lattimore, 1894). Reflecting on depictions of automata in the Iliad, Aristotle speculated in his *Politics* that they might lead to the abolition of slavery by essentially taking over all labor (Aristotle & Ellis, 1888). Another of Hephaestus' creations, the Greek god Talos, is, in a popular telling of the story, forged to protect the goddess Europa who resides on the island of Crete, from pirates and aggressors. Created for one specific purpose, he diligently carries it out never breaking from protocol or upturning the rules (Mayor, 2018).

In a similar fashion to the musical automaton created by Yan Shi in the *Liezi*, Hero of Alexandria (c. AD 62) produced illustrated designs for a range of automata, a good deal of which were musical (Woodcroft, 1851). His *Pneumatics*, building upon and interpreting a text of the same name by Philo of Byzantium (250 BC), provides detailed illustrated designs he produced for theatrical performances. It included an automaton that sounds a trumpet with compressed air, a singing blackbird and a trumpet playing Triton driven by a steam boiler, and two designs for altar organs, one blown by a windmill and another by annual labor (Woodcroft, 1851). His other surviving work *On Automata-Making* contains a range of designs for theatre automata. These were elaborate sets that could involve multiple moving figures, flowing liquids, mechanisms for producing sounds and often some type of fire. Both books were rediscovered during the renaissance and had great influence in the late Renaissance period (Steadman, 2021). Steadman highlights a sense of fun that surrounded automata in the Renaissance period, and this same sense of fun would reemerge from time to time in, for example, Vaucanson's Flute Player, Tambourine Player, and digesting duck automata in the 1700s and Joseph Faber's Fabulous Talking Machine or Euphonia exhibited in 1845 (Riskin, 2003). The earliest designs for truly programmable musical automata can be found in *The Book of Knowledge of Ingenious Mechanical Devices* by 13th-century polymath Ibn Ismail Ibn al-Razzaz Al-Jazari (Sharkey, 2007).

Many of these earlier automata were designed to entertain the rich and powerful with mechanistic analogies of music, dance and theatre. However, during the industrial revolution, the relationship between humans and machines would undergo another transformation. As Western nations began to transition from agriculture and artisan economies to heavy industries, the machines filling the factory floors of the West became critical components of economic infrastructure. Automated textile equipment, in particular, threatened the livelihoods of skilled weavers who could now be replaced with cheaper, less-skilled workers. In 19th-century England, this tension led to the rise of the Luddites, who protested the use of automation to circumvent fair labor practices. They are often remembered as technophobes who feared that they would be replaced by technological innovations. However, writing at the height of Reaganism and Thatcherism, David Linton parallels the economic, social and political conditions that gave rise to the Luddites with the contemporary trends of the day, arguing that their real issue was the increasingly excessive mistreatment and exploitation of labor by capital, of which the industrial machine became a symbol (Linton, 1985). The industrial revolution saw a trend toward wringing more labor from the individual worker at increasingly cheaper rates against a backdrop of continuous decline in both working conditions and living standards. In this landscape, the machine plays the role of diligent servant, the reliable capital of the industrialist class. These machines are not yet the intelligent agents discussed in ancient myth and modern science fiction, yet they have managed to capture and reproduce, in at least some coarse manner, skills and capacities that were previously thought to be uniquely human. This trend has continued into recent times where discussions about the drawbacks of AI in the workplace tend to focus on workers' fears of being replaced as a result of employers' deployment of technology to reduce labor costs and expand the bottom line.

Proponents of the Kurzwellian vision of AI utopia claim that as we approach the dawn of superintelligent AI, an increasing number of professions will be automated, effectively creating a post-labor world where machines do the work and humans reap the rewards, though these rewards tend to involve increased leisure time rather than any share of the profits generated (Skidelsky, 2020). While the old master-servant dynamic is at play here once again, it is also worth noting that superintelligent AIs described by Kurzweil, Vinge and Bostrom are by no means guaranteed to emerge. Alongside those discussed previously (Pein, 2018; Benthall, 2017), a wide range of thinkers, including Tozer (2020) and Colton (2020), argue that AI is radically different from human intelligence and is unlikely to ever operate in a similar or even comparable manner to human cognition.

Promises and Limitations of AI

One early critic of human-level AI was Hubert Dreyfus. During the first wave of AI, Dreyfus argued that symbolically mediated cognitive processes require a context of tacit, informal background knowledge, in the sense indicated by Polanyi (1958), to render them meaningful (Dreyfus, 1965). A large portion of human knowledge, for example, domain-specific expertise, is tacit and informal and so cannot be represented symbolically. Thus, computation alone cannot account for knowledge with a tacit component (Dreyfus, 1976, 1992). The Chinese Room problem (Searle, 1980) showed that while rule-based computation may be sufficient to pass the Turing test, computation alone cannot account for how the symbols computed are assigned their meaning in a human mind. Harnad (1990) formalizes this as the symbol-grounding problem and argues that human experience is full of symbols—the meanings of which computation alone cannot account for. Ragnar Fjelland (2020) highlights that AI research since Dreyfus' original writings has shifted away from the GOFAI model of hard-coding systems of rules toward the design of neural network architectures that can learn relationships directly from a dataset. As such, we may be tempted to believe that these architectures can handle tacit knowledge, but they cannot. This is because the computer does not inhabit the same world as the human does. As such, for example, they do not deal with tacit knowledge as tacit knowledge but rather make arbitrary quantifications (i.e., encode a series of weights), which represent tacit knowledge as explicit knowledge or at least with an explicit formalization. For Fjelland, the DL revolution has not solved this problem because the problem is inherent to computation in general as opposed to AI/ML alone. The flip-side to this argument is that the degree to which we can ground an AI system in Fjelland's human world, mirrors the degree to which it can be said to have a human-like intelligence. But machine intelligence does not need to resemble human intelligence in to be effective as demonstrated in the capabilities of the systems discussed in this chapter. Nor should it need to in order for us to consider it a legitimate form of intelligence.

It may be more useful to think of machine intelligence as characterized less by the kinds of problem solving associated with human cognition and more with the kinds or problem solving we see in the expression of genetic code in the cells and tissues of the biological human substrate. Much of the value of ML/AI lies in the fact that it is different to human intelligence and can therefore do things that humans cannot, just as human intelligence can do things that machines cannot. The computer might not inhabit the human world, as Fjelland notes, but we humans increasingly inhabit the world of computing shaping our lives around it and carrying out much of our daily activities within a landscape of computation. As ML/AI technologies continue to restructure that digital landscape, so too will it restructure our lives and societies. In the end, for Fjelland, Dreyfus' critiques still hold in our current ML-driven epoch of AI. All this being said, the very idea of AI is predicated on the idea that certain machines have the capability to, at the very least, simulate some aspects of human cognition. This assumption of similarity between mind and machine, while present to some degree in the historical and mythological materials explored earlier, began to dominate culture in the 20th century.

The Computer and the Mind

The period during and immediately after the Second World War saw huge advances in the field of computer science. A growing body of research seemed to be lending credence to the idea that the human mind was essentially an information processing machine. The Church-Turing thesis formalized the algorithm (Turing, 1936; Church, 1936) and Turing's theoretical universal machine (1936) claimed to simulate any algorithm with four simple rules. Shannon

(1938) argued that cognitive processes could be modeled by formalizing problem-solving across relay switch states with Boolean logic. McCulloch and Pitts (1943) devised logical models for neural networks and claimed mental activity thus could be modeled on a universal Turing machine. Von Neumann (1945) made it possible for a machine to program itself by storing programs in memory. Finally, at the first Hixon Symposium on Cerebral Mechanisms and Behaviour in 1948, thinkers like John Von Neumann, Warren McCulloch and Walter Pitts addressed talks to a multidisciplinary crowd of psychologists and computer scientists, in which they approached the brain, central nervous system and the mind in terms of computation (Gardner, 1987). This incitement, coupled with a growing sense of dissatisfaction with the methods and results of the prevailing behaviorist agenda in psychology, would contribute to the establishment of cognitive science as a distinct field of research. A key assumption on which the field was founded was the equivalence between mind and machine and the treatment of the human as an information processor. When AI emerged as a research field in the 1950s this model of cognition, what Putnam (1967) would later dub the classical computational theory of mind, was well established. Many of the key players in the early days of both cognitive science and computer science would also play roles in the development of the field. This line of thought prevailed in cognitive science until the 1980s when the impact of pioneering work in the field of cybernetics would be felt in cognitive science too. In fact, Hayles (1999) describes how cybernetics wrestled with this interpretation of the human, almost from its inception.

Cybernetics and Cybernetic Art

Cybernetics originated with the work of mathematician Norbert Wiener who in his seminal 1948 text described it as the scientific study of control and communication in the animal and the machine (Wiener, 1948). He saw goal-directed, teleological behaviors such as self-regulation through feedback and feedforward loops as fundamental to electronic, mechanical and biological systems. Second-order cybernetics built on these ideas and driven by the work of Margaret Mead (1968) and Heinz Von Foerster (1984) introduced reflexive practices whereby the observing agent became a critical feature of the system. Maturana and Varela's autopoiesis (1980) and Stafford Beer's eventual reconfiguration of cybernetics from an operations research perspective as "the science of effective organization" (1972) would further expand the cybernetic horizon.

Cybernetic thought had a great influence on the arts, as demonstrated when in 1968 Jasia Reichardt curated an exhibition at the ICA in London titled Cybernetic Serendipity (Reichardt, 1968), which explored the relationship between technology and creativity. The exhibition was a watershed moment in new media arts. It brought together a wide range of artists, engineers, mathematicians and architects working within a cybernetic framework. It included contributions from, among others, Gordon Pask, Stafford Beer, Jeanne

Beaman, Frieder Nake, J. R. Pierce, Peter Zinovieff, Gustav Metzger, Nam June Paik, Frank Malina, Roger Dainton, John Cage, Karlheinz Stockhausen, Maughan S. Mason, A. R. Forrest and Margaret Masterman.

In the book accompanying the exhibition, Reichardt comments that "it may be difficult for an artist to imagine how he could possibly make use of a computer. The solution to the problem lies in collaboration" (Reichardt, 1968). This theme of collaboration between artist and machine is extended to the audience in Pask's *Colloquy of Mobiles* (Pask, 1969). Pask was a prominent cybernetic theorist with a deep interest in interactive installation art. His was an "aesthetically potent" environment that allowed its audience to actively engage in a discourse with a hierarchy of interacting mobiles. He favored a collaborative relationship between machine and human agents, a theme that was also present in his other works *Musicolour* and *SAKI* (Bird & DiPaolo, 2008).

Hayles (1999) points out that Cybernetics came together as a distinct field of research during the Macy conferences of 1946 to 1953 and in its initial wave was defined by a dialectic tension between the homeostasis, the ability of a system to maintain itself in a stable state through corrective feedback, and reflexivity, the tendency of a system to evolve, change and complexify in response to self-observation. Hayles highlights how the homeostasis camps were influenced by Claude Shannon's application of his information theory to cybernetics. They viewed the human as an information processing machine made noisy and erratic by the psychological complexity of subjective experience, and which may in time be replaced by a more efficient information processing machine. The reflexivity camp meanwhile felt that the subjectivity and psychological complexity of the human were to be embraced and accounted for in the negotiation of a more open and collaborative relationship between human and machine. Donald M. MacKay in particular viewed reflexivity as a reconciling agent between information and meaning, the latter of which is wholly unaccounted for in Shannon's information theory. The model of the human as an information processing machine, and an unreliable one at that, must inevitably lead to the conclusion that the human machine is fated for eventual replacement by a more efficient information processing machine. This theme reemerges again today in the cultural discourse around machine intelligence.

While Cybernetic Serendipity was generally well received, Usselmann (2003) reviews some of the dissenting voices at the time including some who rejected what they saw as a kind of techno-fetishist project to replace artists with machines and others who felt the event represented a kind of technocratic authoritarianism. Writing in 1971 Reichardt (Reichardt, 1971) criticized the tendency among journalists to ask if computers would replace humans, arguing that this emotionally charged line of questioning serves to obscure rather than illuminate the relationship between art and technology, further adding that the demystification of the art-making process does not demystify the result. For Reichardt a cybernetic approach to the arts is about human-machine collaboration as opposed to replacement of the artist or the use of machines by one class of humans to oppress another. To this degree, her view of the role of the human is closer to MacKay's than Shannon's and her collaborative take on cybernetics in the arts addresses both the master-servant dynamic and replacement anxiety.

While much of the work presented at Cybernetic Serendipity engaged with concepts and ideas from the field, the ten-track album released alongside the show, titled Cybernetic Serendipity Music, was also representative of the state of the art in experimental music as much as it represented human-machine collaboration. Alongside works by Zinovieff, Brün and Strang it included the fourth movement of Hiller and Isaacson's *Illiac Suite* (Hiller & Isaacson, 1957), an excerpt of Xenakis' *Stratégie* (1962) and an excerpt of Cage's *Cartridge Music* (Various Artists, 1968). Each of the pieces generally involved the use of a computer to generate musical materials through either the application of some ruleset or algorithm-like game theory on Xenakis' *Stratégie*, compositional processes hand-coded in FORTRAN on Herbert Brün's *Infraudibles*, or stochastic processes in Cage's *Cartridge Music*.

Perspectives on Experimental Music

Boden (2004) argues that technologically mediated approaches in the creative arts necessarily involve the demystification of the creative process in explicit systems, models and rules. Furthermore, such approaches invariably involve the ceding of agency from artist to machine to some in regard. In the cases where the systems are overly constrained in terms of what outputs they can produce, creative works can become predictable and dull. This dialectic between what is boring and predictable and what is exciting touches all art but technologically mediated art especially (Nake, 2012). This becomes especially problematic in the current DL era, where, in a certain sense, predictability is fundamentally baked into the architecture of any algorithm. While ML models may generate new samples not present in their original input data, the set of all possible sample predictions is nonetheless defined and constrained by relationships and patterns in the original data, as encoded at training time by the architecture of the model. This throws up some interesting problems in the context of experimental music in particular. Cage (1961) held that, in the context of music composition, "an experimental action is one the outcome of which is not foreseen" and is "necessarily unique." It is not possible to take a body of known and well-understood musical data, and to create, by a process of prediction, an unforeseen musical output, which satisfies Cage's definition for experimental music. Furthermore, this problem cannot be solved by simply training a generative model on a corpus of experimental musical works. Important examples of musical works at the more experimental end of the spectrum tend to be highly specific one-of-a-kind pieces with unique reasons for existing and completely idiosyncratic aesthetics and internal organizational schemes. It is hard to argue that those patterns that are shared between Hugh Le Caine's *Dripsody* (1955) and Steve Reich's *It's Gonna Rain* (1965) are definitive of "experimental music" as an art form or similarly that any set of correlations between Xenakis' *Metastasis* (1955) and Halim El-Dabh's *Wire Recorder Piece* (1944) could be elaborated upon to generate a new experimental music composition. This is to say nothing of installed works, such as La Monte Young and Marian Zazeela's *Dream House* (1969) or Gordon Monahan's *Aeolian Silo* (1990), that cease to operate in the same manner when divorced from their context. These are all but unrepresentable in audio or video file formats and so it would prove exceptionally difficult to represent the original pieces in an ML dataset.

The patterns that exist across pieces of experimental music are much more loosely correlated than those that exist across more heavily regimented forms of musical expression like Western classical music and pop. The measure of distance between the familiar examples of experimental music is too large. Each one differs wildly from the next in terms of texture, timbre and form. Statistically speaking, a dataset comprising important or familiar experimental pieces would be inherently "noisy," in the sense described by Shannon's information theory, as there are few common patterns across pieces. Continuing with this metaphor, those patterns in the dataset that are stable and repetitive enough across examples to constitute a "signal" are fundamentally at odds with the spirit and intention of experimental music. Experimental music, in the sense described by Cage, cannot be produced as a prediction from a model trained on a musical dataset but could only result from a complex system or chaotic system in which the musical result cannot be foreseen. The idea of a musical result generated as a prediction from a model determined by other musical works is at odds with the Cagean understanding of experimental music, which favors randomness over predictability.

However, Cage's definition of experimental music, while historically privileged among Western Eurocentric scholars and artists, is not the only definition of experimental music. Critically, it is not a definition of, or approach to, experimental music that resonates within the current epoch of AI-driven art. Cage's musical ideas were influenced by his interest in Daisetz Suzuki's particular expression of Zen Buddhism, which was heavily influenced by William James' concept of "Pure Experience" as mapped to Zen in the early work of Nishida Kitarō (Roddy, 2017). George E. Lewis (1996) points out that Cage's approach, in its focus on unique and spontaneous chance operations in the present moment, eliminates personality, narrative, memory and history. In doing so it minimizes the works of prominent African American Jazz composers and improvisers like Charlie Parker, Dizzy Gillespie and Thelonious Monk who produced truly experimental music that embraced these elements in a live group improvisation context. He further highlights how Cage outright dismissed the value of Jazz in demarcating it from what he called "serious music" even when, as Born (1995) points out, Jazz musicians had incorporated core elements of experimental music practice since the 1950s. Lewis's take on experimental music is deeply indebted to Charlie Parker's idea of music as an expression of the lived experience of the players involved. Lewis sees collaborative musical improvisation steeped in personal narrative as an assertion of agency on behalf of the musicians involved. This is in stark contrast to Cage and followers whose master-servant model of composition called for the enforcement of control over the musical performer. This thinking underlies Lewis's development of Voyager, an interactive computer music environment that operates as a virtual improvising orchestra (Lewis, 2000) that was initially developed by combining principles from 1980s AI, and 1950s cybernetics, with sociomusical networks of free improvisation (Lewis, 2019, 2021). Debuted in 1987, Voyager would develop and evolve over the years, always operating as an autonomous computer-based system capable of improvising intelligently alongside human performers in a live context (Steinbeck, 2018). It is still an ongoing project and in 2022 Lewis began working with a team of researchers and practitioners at RNCM PRiSM to expand Voyager integrating AI/ML techniques (Royal Northern College of Music, 2022). For Lewis, collaborative improvisation is an assertion of agency that breaks from the master-servant dynamic, which pervades both human-machine collaboration and the relationship between composer and performer as reimagined by Cage and his followers. This contrasts with earlier musical machines like the automatons, which were designed to entertain rich and powerful masters, by reproducing set musical pieces to the rigid specification of the original designer/composer. Lewis's approach presents a model of experimental music that grants agency and autonomy to musical machines as improvisers of equal value alongside their flesh and blood collaborators.

Current Trends and Future Directions

Lewis's approach to the use of intelligent machines as equal collaborators addresses both the issue of the master-servant dynamic and the fear that humans will be replaced by machines. It acknowledges and builds upon the reflexive approach of those second-order cyberneticists who believed human observers, with their complex histories, memories, personalities and individual autonomy constituted critical components of cybernetic systems. It takes this a step further however by extending the same courtesy to the machine and treating machine agents as equal collaborators that act with autonomy alongside their human counterparts. It further accounts for the cultural and social context in which these machine-human collaborators are performing, representing a highly novel and even compassionate approach to the use of intelligent machines in the creative arts. There are shades of Lewis's thinking emerging in the current epoch of AI/ML-driven art. Sougwen Chung's work explores equitable machine-human collaboration. Her piece *Omnia per Omnia* (Chung, 2018) approaches landscape painting as a collaboration between artist, robotic swarm and city. Špela Petrič's *PL'AI* (Petrič, 2020) allowed an AI-driven robot and a set of plants to interact with each other over a long timescale through play. The author's *Signal to Noise Loops* project foregrounds collaboration between city, machine and human in electronic music and audiovisual installation contexts (Roddy, 2023). The emergence of projects of this type that carry on the tradition of the cybernetics represented in George E. Lewis's work, free from the master-servant dynamic and anxiety of replacement, bode well for the future of collaboration between machines and humans in the creative arts.

References

Akten, M. (2017). *Hello world* [interactive art]. https://www.memo.tv/works/learning-to-see/

Anadol, R. (2016). *Machine hallucinations* [visual art]. https://refikanadol.com/works/machine-hallucinations-nature-dreams/

Aristotle, & Ellis, W. (1888). *Politics: A treatise on government from the Greek of Aristotle*. Routledge.

Audry, S. (2021). *Art in the age of machine learning*. MIT Press.

Beer, S. (1972). *Brain of the firm*. Wiley.

Benthall, S. (2017). *Don't fear the reaper: Refuting Bostrom's superintelligence argument*. arXiv preprint arXiv:1702.08495. https://arxiv.org/abs/1702.08495

Bird, J., & Di Paolo, E. (2008). Gordon Pask and his maverick machines. *The mechanical mind in history*, 185–211.

Boden, M. A. (2004). *The creative mind: Myths and mechanisms*. Routledge.

Born, G. (1995). *Rationalizing culture: IRCAM, Boulez, and the institutionalization of the musical Avant-Garde*. University of California Press.

Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies* (1st ed.). Oxford University Press.

Breuer, H. P. (1975). Samuel Butler's "the book of the machines" and the argument from design. *Modern Philology*, *72*(4), 365–383.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., . . . Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, *33*, 1877–1901.

Butler, S. (2015). *Erewhon and erewhon revisited*. Courier Dover Publications.

Cage, J. (1961). *Silence: Lectures and writings*. Wesleyan University Press.

Cardew, B. (2018, September 24). Ash Koosha: Return 0. *Pitchfork*. Retrieved August 5, 2022, from https://pitchfork.com/reviews/albums/ash-koosha-return-0/

Chung, S. (2018). Omnia per Omnia (2018)—Sougwen Chung (愫君). *Sougwen.com*. Retrieved August 9, 2022, from https://sougwen.com/project/omniaperomnia

Church, A. (1936). A note on the entscheidungsproblem. *The Journal of Symbolic Logic*, *1*(1), 40–41.

Cohen, H. (1982, December). How to make a drawing. In *Talk given to the science colloquium* (Vol. 17). National Bureau of Standards.

Colton, S. (2020). Possibilities and limitations for AI: What can't machines do? In R. Skidelsky & N. Craig (Eds.), *Work in the future: The automation revolution*. Springer Nature.

Cope, D. (1992). Computer modeling of musical intelligence in EMI. *Computer Music Journal*, *16*(2), 69–83.

Darville, J. (2020, December 16). Arca shares 100 different remixes of "Riquiquí" generated by A.I. *The Fader*. Retrieved August 5, 2022, from www.thefader.com/2020/12/16/arca-shares-100-different-remixes-of-riquiqu-generated-by-ai

Dhariwal, P., Jun, H., Payne, C., Kim, J. W., Radford, A., & Sutskever, I. (2020). Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*.

Dhariwal, P., & Nichol, A. (2021). Diffusion models beat GANs on image synthesis. *Advances in Neural Information Processing Systems*, *34*, 8780–8794.

Dick, P. F. (1968). *Do androids dream of electric sleep?* Doubleday.

Dreyfus, H. L. (1965). *Alchemy and artificial intelligence*. RAND Corporation.

Dreyfus, H. L. (1976). What computers can't do. *British Journal for the Philosophy of Science*, *27*(2).

Dreyfus, H. L. (1992). *What computers still can't do: A critique of artificial reason*. MIT Press.

Eck, D., & Schmidhuber, J. (2002). *Finding temporal structure in music: Blues improvisation with LSTM recurrent networks*. IEEE Workshop on Neural Networks for Signal Processing.

El-Dabh, H. (1944). *Wire recorder piece* [composition]. https://classical20.com/2013/01/15/halim-el-dabh-wire-recorder-piece-1944/

Elgammal, A. (2019). AI is blurring the definition of artist: Advanced algorithms are using machine learning to create art autonomously. *American Scientist*, *107*(1), 18–22.

Ellison, H. (1967). I have no mouth, and i must scream. In *IF: Worlds of science fiction* (pp. 162–175). Galaxy Publishing Corp.

Engel, J., Hantrakul, L., Gu, C., & Roberts, A. (2020). *DDSP: Differentiable digital signal processing*. arXiv preprint arXiv:2001.04643. https://arxiv.org/abs/2001.04643

Esser, P., Rombach, R., & Ommer, B. (2021). Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 12873–12883). IEEE, CVF.

Fiebrink, R., Trueman, D., & Cook, P. R. (2009). *A meta-instrument for interactive, on-the-fly machine learning*. https://soundlab.cs.princeton.edu/publications/FiebrinkTruemanCook_NIME2009.pdf

Fjelland, R. (2020). Why general artificial intelligence will not be realized. *Humanities and Social Sciences Communications*, *7*(1), 1–9.

Flynn, M. (2018). A 19-year-old developed the code for the AI portrait that sold for $432,000 at Christie's. *Washington Post*, *26*.

Fry, A. (2021). EDM, meet AI. *The Evolution of Electronic Dance Music*, 81–102.

Gardner, H. (1987). *The mind's new science: A history of the cognitive revolution*. Basic Books.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, *27*.

Google. (2022a, May). *Imagen*. Retrieved August 4, 2022, from https://imagen.research.google/

Google. (2022b, June). *Parti*. Retrieved August 4, 2022, from https://parti.research.google/

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, *42*(1–3), 335–346.

Hayles, K. N. (1999). *How we became Posthuman: Virtual bodies in cybernetics, literature, and informatics* (1st ed.). University of Chicago Press.

AuQ2 Heikkilä, M. (2023, March 21). How AI experts are using GPT-4. *MIT Technology Review*. https://www.technologyreview.com/2023/03/21/1070102/how-ai-experts-are-using-gpt-4/

Herndon, H. (2021, July 13). *Holly+*. Retrieved August 5, 2022, from https://holly.mirror.xyz/54ds2IiOnvthjGFkokFCoaI4EabytH9xjAYy1irHy94

Hiller, L. A., & Isaacson, L. M. (1957, October). Musical composition with a high speed digital computer. In *Audio engineering society convention 9*. Audio Engineering Society.

Jancer, M. (2018, May 17). More artists are writing songs in the key of AI. *Wired*. Retrieved August 5, 2022, from www.wired.com/story/music-written-by-artificial-intelligence/

Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., & Aila, T. (2021). Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, *34*, 852–863.

Kirn, P. (2020, July 6). Transfiguración: decolonizing AI, in Hexorcismos' shamanistic music and art - CDM Create Digital Music. *CDM Create Digital Music*. https://cdm.link/2020/07/transfiguracion-decolonizing-ai-in-hexorcismos-shamanistic-music-and-art/

Klingemann, M. (2017). *Imposture series* [visual art]. https://www.artsy.net/artwork/mario-klingemann-imposture-series-cobalamime

Kuivila, R. (2004). David Tutor: Live electronic music. *Leonardo Music Journal*, *14*, 106–107.

Kurzweil, R. (2005). *The singularity is near: When humans transcend biology*. Penguin.

Lattimore, R. (Ed.). (1894). *The Iliad of Homer*. Cambridge University Press.

Le Caine, H. (1955). *Dripsody: An étude for variable speed recorder* [composition]. NRC.

Lem, S. (2021). Lymphater's formula. In A. Lloyd-Jones (Ed.), *The truth and other stories* (pp. 221–248). MIT Press.

Lewis, G. E. (1996). Improvised music after 1950: Afrological and eurological perspectives. *Black Music Research Journal*, 91–122.

Lewis, G. E. (2000). Too many notes: Computers, complexity and culture in voyager. *Leonardo Music Journal*, *10*, 33–39.

Lewis, G. E. (2019). Listening for freedom with Arnold Davidson. *Critical Inquiry*, *45*(2), 434–447.

Lewis, G. E. (2021). Co-creation: Early steps and future prospects. In Bernard Lubat, Gérard Assayag, Marc Chemillier (Eds). Artisticiel / Cyber-Improvisations. Phonofaune, 2021, Dialogiques d'Uzeste., 2021. hal-03543133.

Lewis, J. P. (1988, July). Creation by refinement: A creativity paradigm for gradient descent learning networks. In *International conference on neural networks* (pp. 229–233). IEEE.

Linton, D. (1985). Luddism reconsidered. *ETC: A Review of General Semantics*, 32–36.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living (Boston studies in the philosophy and history of science, 42)* (1980th ed.). Springer.

Mayor, A. (2018). Gods and robots. In *Gods and robots*. Princeton University Press.

McCorduck, P. (1979). Machines Who Think: A Personal Inquiry into the History and Prospects of. *Intelligence*. San Francisco: Freeman.

McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, *5*(4), 115–133.

Mead, M. (1968). The cybernetics of cybernetics. In H. von Foerster, J. D. White, L. J. Peterson, & J. K. Russell (Eds.), *Purposive systems* (pp. 1–11). Spartan Books.

Midjourney. (2023, Mar). *Midjourney*. https://www.midjourney.com. Retrieved March 23, 2023, from https://www.midjourney.com

Miller, A. I. (2019). *The artist in the machine: The world of AI-powered creativity*. MIT Press.

Monahan, G. (1990). *Aeolian silo* [installation]. Funny Farm.

Nake, F. (2012). Construction and intuition: Creativity in early computer art. In *Computers and creativity* (pp. 61–94). Springer.

AuQ3 Oord, A. V. D., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., & Kavukcuoglu, K. (2016). *WaveNet: A generative model for raw audio* (p. 125). Proceedings 9th ISCA Workshop on Speech Synthesis Workshop (SSW 9).

OpenAI. (2021a, January). Dall-E 2. *OpenAI.com*. Retrieved August 4, 2022, from https://openai.com/dall-e-2/

OpenAI. (2021b, January). CLIP: Connecting text and images. *Openai.com*. Retrieved August 4, 2022, from https://openai.com/blog/clip/

Park, T., Liu, M. Y., Wang, T. C., & Zhu, J. Y. (2019). GauGAN: Semantic image synthesis with spatially adaptive normalization. In *ACM siggraph 2019 real-time live!* (p. 1). ACM.

Pask, G. (1969). *Colloquy of mobiles* [interactive art]. The ICA.

Payne, C. M. (2021, June 21). MuseNet. *OpenAI.com*. Retrieved August 4, 2022, from https://openai.com/blog/musenet/

Pein, C. (2018). *Live work work work die: A journey into the savage heart of silicon valley*. Metropolitan Books.

Petrič, Š. (2020). PL'AI. *Špela Petrič.org*. Retrieved August 8, 2022, from www.spelapetric.org/#/plai/

Polanyi, M. (1958). *Personal knowledge: Towards a post-critical philosophy* (1st ed.). University of Chicago Press.

Putnam, H. (1967). Psychological predicates. *Art, Mind, and Religion*, *1*, 37–48.

Quarshie, A. (2021, September 4). The quietus | features | a quietus interview | technology rewires your brain: Lee Gamble interviewed. *The Quietus*. Retrieved August 5, 2022, from https://thequietus.com/articles/30510-lee-gamble-interview

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., & Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, *21*(140), 1–67.

Reich, S. (1965). *It's gonna rain* [composition]. https://en.wikipedia.org/wiki/It%27s_Gonna_Rain

Reichardt, J. (1968). *Cybernetic serendipity: The computer and the arts* (1st ed.). Studio International.

Reichardt, J. (1971). *Cybernetics, art and ideas*. Studio Vista.

Richey, J. L. (2011). I, robot. *Riding the Wind with Liezi: New Perspectives on the Daoist Classic*, *193*.

Riskin, J. (2003). The defecating duck, or, the ambiguous origins of artificial life. *Critical inquiry*, *29*(4), 599–633.

Roberts, A. (2016). *The history of science fiction*. Palgrave Macmillan.

Roberts, A., Engel, J., Raffel, C., Hawthorne, C., & Eck, D. (2018, July). A hierarchical latent vector model for learning long-term structure in music. In *International conference on machine learning* (pp. 4364–4373). PMLR.

Roddy, S. (2017). Absolute nothingness—the Kyoto school and sound art practice. In R. Castro (Ed.), *Proceedings of the 2017 invisible places conference* (pp. 156–172). Invisible Places.

Roddy, S. (2023). Signal to noise loops: A cybernetic approach to musical performance with smart city data and generative music techniques. *Leonardo*, *56*(1).

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 10684–10695). IEEE, CVF.

Royal Northern College of Music. (2022, March 18). *George Lewis residency | RNCM musicians and an enhanced voyager—Royal Northern college of music*. RNCM. Retrieved

August 7, 2022, from www.rncm.ac.uk/research/research-centres-rncm/prism/prism-news/george-lewis-residency-rncm-musicians-and-an-enhanced-voyager/

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*(3), 417–424.

Shanahan, M. (2015). *The technological singularity*. MIT Press.

Shannon, C. E. (1938). A symbolic analysis of relay and switching circuits. *Electrical Engineering*, *57*(12), 713–723.

Sharkey, N. (2007). I ropebot. *New Scientist*, *195*(2611), 32–35.

Sherburne, P. (2021, March 1). Mouse on mars: AAI. *Pitchfork*. Retrieved August 5, 2022, from https://pitchfork.com/reviews/albums/mouse-on-mars-aai/

Skidelsky, R. (2020). The future of work. In R. Skidelsky & N. Craig (Eds.), *Work in the future: The automation revolution*. Springer Nature.

SKYGGE. (2017). *Hello world album the first album composed with an artificial intelligence*. Retrieved August 4, 2022, from www.helloworldalbum.net/

Steadman, P. (2021). The automata of Hero of Alexandria. *Renaissance Fun: The Machines Behind the Scenes*, 111–132.

Steinbeck, P. (2018). George Lewis's voyager. In *The Routledge companion to jazz studies* (pp. 261–270). Routledge.

Todd, P. (1988). A sequential network design for musical applications. *Proceedings of the 1988 Connectionist Models Summer School*, 76–84.

Toffler, A. (1970). *Future shock*. Random House.

Tozer, T. (2020). What computers will never be able to do. In R. Skidelsky & N. Craig (Eds.), *Work in the future: The automation revolution*. Springer Nature.

Tudor, D. (1995). *Liner notes* [CD]. Lovely Music.

Turing, A. (1936). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, *42*(1).

Tyka, M. (2018). *Portraits of imaginary people* [visual art]. https://www.interaliamag.org/audiovisual/mike-tyka/

Usselmann, R. (2003). The dilemma of media art: Cybernetic serendipity at the ICA London. *Leonardo*, *36*(5), 389–396.

Various Artists. (1968). *Cybernetic serendipity music* [record]. Institute of Contemporary Arts.

Vinge, V. (1993). The coming technological singularity. *Whole Earth Review*, *81*, 88–95.

von Foerster, H. (1984). *Observing systems*. Intersystems Publications.

von Neumann, J. (1945, June 30). *First draft of a report on the EDVAC* [Contract No. W]. https://en.wikipedia.org/wiki/First_Draft_of_a_Report_on_the_EDVAC

Warzel, C. (2023, February 8). Talking to AI might be the most important skill of this century. *The Atlantic*. https://www.theatlantic.com/technology/archive/2023/02/openai-text models-google-search-engine-bard-chatbot-chatgpt-prompt-writing/672991/

Wiener, N. (1948). *Cybernetics or control and communication in the animal and the machine*. MIT Press.

Woodcroft, B. (Ed.). (1851). *The pneumatics of Hero of Alexandria: From the original Greek*. Taylor, Walton and Maberly.

Xenakis, I. (1955). *Metastasis* [musical performance]. Donaueschingen Festival.

Xenakis, I. (1962). *Stratégie*. London: Boosey & Hawkes.

Young, L., & Zazeela, M. (1969). *Dream house* [installation]. https://en.wikipedia.org/wiki/Dream_House_(installation)

Zukowski, Z., & Carr, C. J. (2018). *Generating black metal and math rock: Beyond Bach, Beethoven, and Beatles*. arXiv preprint arXiv:1811.06639. https://scholar.google.com/citations?user=P7dSScMAAAAJ&hl=en