

Deep Listening: Early Computational Composition and its Influence on Algorithmic Aesthetics

Tiffany Funk
University of Illinois at Chicago
Chicago, Illinois 60607
tfunk4@uic.edu

The popularity of the program DeepDream, created by Google engineer Alexander Mordvintsev, exploded due to the open source availability of its uncanny, hallucinogenic aesthetic. Once used to synthesize visual textures, the program popularized the concept of neural network training through image classification algorithms, inspiring visual art interrogating machine learning and the training of proprietary prediction algorithms; though DeepDream has facilitated the production of many mundane examples of surreal computer art, it has also helped to produce some conceptually rich visual investigations, including MacArthur “genius” awardee Trevor Paglen’s recent installation *A Study of Invisible Images*. While the significance of trained neural networks is presently considered valuable to computer vision experimentation, a medial archaeological investigation of the conceptual underpinnings of machine learning reveals the fundamental influence early *sonic* experiments in computational music have in its computational and conceptual framework. Early computational music works, such as Lejaren Hiller Jr. and Leonard Isaacson’s *Illiad Suite* (1957), the first score composed by a computer, as well as Hiller and John Cage’s ambitious multimedia performance *HPSCHD* (1969), used stochastic models to automate game-like processes, such as Giovanni Pierluigi da Palestrina’s Renaissance-era polyphonic instruction, as well as the *I Ching* divination process of casting coins or yarrow stalks. Hiller’s concerns regarding the historical use of compositional/mathematical gameplay uncovers a conceptual and performative emphasis anticipating the “training” of visual models. Through the adverse reactions of audiences to Hiller’s compositions, written by what the press deemed derogatorily “An electronic brain” in 1957 parallel public reactions to the disturbing mutations of DeepDream, popular participation in the open source project signals a growing willingness to collaborate creatively with computers to interrogate both computational and cognitive processes.

Machine learning; music composition; aesthetics.

1. INTRODUCTION

Just as an oft-quoted quip by Voltaire dictates, “The Holy Roman Empire was neither Holy, Roman, nor an Empire,” one might say the same thing about machine learning (ML): though it is a crucial branch of current Artificial Intelligence (AI) research, ML does not actually involve “learning,” but rather projecting with the use of big data. In fact, in more techno-skeptical circles, ML is just a trendier definition for linear algebraic systems too boring and time-consuming for humans to parse. Despite the fatigue and skepticism surrounding the intertwining fields of ML, deep learning, and pattern recognition research, it’s hard to deny the eerie, yet seductive qualities of images constructed by trained, image classification neural networks; for example, the deliberate over-processed computer vision images from DeepDream, created by Google engineer Alexander Mordvintsev, produced an international stir on social media with the open source availability of its uncanny, psychobolic aesthetic (see Figure 1).

However, current focus on computer vision experimentation neglects its sonic precursors: a reexamination of the institutional, algorithmic

research in computational music composition during the 1950s and 60s reveals a compelling historical pretext to contemporary machine learning experimentation. [1] These early stochastic music compositions, painstakingly programmed on supercomputers ill-fitted for musical experimentation, were the result of two converging post-World War II phenomena: the first involved the sudden exuberance toward interdisciplinary art and technology experiments, an aftereffect of wartime scientific guilt. Many of these collaborations were the result of the interdisciplinary nature of computing centers within research institutions before the onset of recent disciplinary rigidity.



Figure 1: A late-stage DeepDream processed photograph of three men in a pool. *Deep_Dreamscope_(19822170718).jpg*, by jessica

mullen from austin, tx, Deep Dreamscope, is licensed under CC BY 2.0.

The second postwar phenomenon affecting machine learning's sonic pre-history involved sheer computational limitations: as post-war computers lacked the means to process or store vast quantities of data, speculative arts research in transcribing music for predictive analysis provided one of the few opportunities to craft algorithms pre-dating the training models of machine learning. Lejaren A. Hiller Jr. and Leonard Isaacson's 1957 composition *The Illiac Suite*, often credited as the first score composed by a computer, was the first of a series of early experimental compositions aimed at exploring music history through computational predictive analysis. Assisted by Isaacson, Hiller developed stochastic models with a conceptual and performative emphasis that ultimately exploring computational and human sensory limitations. The contemporary exploitation of these limitations is exemplified by the over-processing of DeepDream's "training" of visual models; both visual and sonic aesthetics are ultimately shaped by the ability to hear tonal complexity or see (and comprehend) the construction of visual patterns organized through statistical probabilistic models.

However, both the sensory and technological limitations facing Hiller reflects larger movements in conceptual art and atonal and serial music composition in the mid 20th century. Hiller's *Illiac Suite*, as an attempt to translate statistical analysis to aesthetic analysis, may be poetically equated to Dutch Baroque painter Johannes Vermeer's exploitation of the camera obscura "disks of confusion." Current theories claim Vermeer's stylistic flourish of *pointillés*, the tiny dabs of pigment representing light refraction, mimic the imperfections of the camera obscura's lens (Mills, 1998). Just as the blurred, unfocused spots created by early optical technologies define Vermeer's style, Hiller's own musical *pointillés* reflect both his computational limitations—the relatively limited processing power of the *ILLIAC* supercomputer in combination with his relatively small data set and meticulously crafted probability tables—as well as his own avant-garde musical influences, primarily Arnold Schoenberg and Milton Babbitt. While his methods pre-date machine learning processes, Hiller's approach to constructing the *Suite*'s complex probability tables anticipated contemporary model-based machine learning practices in service of building a composition exploring meta-historical music aesthetics through computational means. The resulting aesthetic tested the limitations of both computer and human senses, anticipating contemporary visual and sonic experiments that exploit these limitations to ascertain the fluctuations in the gulf between computational "learning" and human intelligence, and what we gain or sacrifice in narrowing that gap.

2. MARKOV CHAIN MONTE CARLO: NASCENT MACHINE LEARNING

Computer scientist Tom Mitchell provides the most concise definition of ML algorithms, stating, "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E" (Mitchell, 1997, p. 2). Notably, Mitchell's definition concerns operation and performance rather than concepts of "thinking" or "learning," following Alan Turing's proposal in his seminal paper "Computing Machinery and Intelligence": he substitutes the question "Can machines think?" with "Can machines do what we (as thinking entities) can do?" (Turing, 1950, p. 433).

Though ML and its related fields—from constructing neural networks to computer vision training—dominates current computational research, the probabilistic methods that define the fields existed long before their popularization. Stanley Ulam, inventor of the statistical sampling technique Markov chain Monte Carlo Method (MCMC), recalled how his method—stringing randomly generated variables representing the present state to model how current changes affect future states—was the result of a boredom-induced round of hospital solitaire:

The first thoughts and attempts I made to practice [the Monte Carlo method] were suggested by a question which occurred to me in 1946 as I was convalescing from an illness and playing solitaires. The question was what are the chances that a Canfield solitaire laid out with 52 cards will come out successfully? ... I wondered whether a more practical method than 'abstract thinking' might not be to lay it out say one hundred times and simply observe and count the number of successful plays. This was already possible to envisage with the beginning of the new era of fast computers, and I immediately thought of problems of neutron diffusion and other questions of mathematical physics, and more generally how to change processes described by certain differential equations into an equivalent form interpretable as a succession of random operations (Eckhardt, 1987, 131).

Ulam's idea was enthusiastically adopted by John von Neumann for applications to neutron diffusion, the name "Monte Carlo" suggested by Nicholas Metropolis. (Eckhardt describes the early history of MCMC, and Hitchcock, 2003, constructs a brief history of the Metropolis algorithm.) Von Neumann *et. al* recognized that MCMC's strength was its ability to construct models based on a sampling from a probability distribution; in cases where there are hundreds or thousands of unknown parameters, MCMC could allow for digital

computation of large hierarchical models that may be adapted over time.

Despite von Neumann's excitement regarding its usefulness, the realization that MCMC and its various subset methods could be used in a wide variety of situations came over forty years later, as recounted by statisticians Christian Robert and George Casella:

In the 1982 edition of the Encyclopedia of Statistical Sciences, there was no entry for "Metropolis-Hastings algorithm," "Metropolis," "Hastings," or "Markov chain Monte Carlo" (Kotz and Johnson 1982)... J. M. Hammersley and D. C. Handscomb, in their 1964 book *Monte Carlo Methods*, mentioned the Metropolis method, but they seemingly fail to grasp its great potential. They listed it as a method of solving "problems in equilibrium statistical mechanics"... [and] not as a general way to simulate observations from virtually any distribution (Hammersley and Handscomb, 1964, pp. 117-121; Robert and Casella, 2011, pp. 245-255).

The relative obscurity of MCMC is thought to be largely due to lack of the computing machinery that made advancements in statistical analysis and ML possible. However, as demonstrated by Hiller and Isaacson's experience programming the *Illiac Suite*, it is also clear that while MCMC necessitated larger data sets and more computing power than was readily available at the time, Hiller's expertise and position, coupled with his dogged drive to test the creative methods he learned as a music student, provided the exact conditions to make the *Suite*, and Hiller's impressive computational music career, possible.

3. THE ILLIAC SUITE AND AI

In the evening of August 9th, 1956, University of Illinois Chemistry Department researchers Lejaren A. Hiller Jr. and Leonard Isaacson debuted the first three movements of the *Illiac Suite: String Quartet No 4* in the Woodward Lounge of the Illini Union (see Figure 2). Though scant documentation of the premiere still exists, secondary sources—for example, the following day's United Press news release (Figure 3)—describes its "resentful" crowd; the release declares that the *Suite*, "COMPOSED BY AN ELECTRONIC BRAIN" and only "SPONSORED BY L.A. HILLER, A CHEMIST-COMPOSER, AND L. M. ISAACSON, A RESEARCH ASSOCIATE," left the self-described "MUSIC LOVER" in a glum state, lamenting it as a death knell for human creativity (Bewley, 2004, p. 11). Though Hiller later dismissed the hyperbole as "rather silly" (1981, p. 10), the *Suite*'s implementation of algorithmic rules describing the history of composition was overshadowed by its novel AI research performed on its namesake, the

Illinois Automatic Computer (or, ILLIAC), the first high-speed supercomputing center in a university (see Figure 4).

Hiller described the *Suite*—what he called "a bootleg job at night" during his days employed as a researcher in the University of Illinois Chemistry Department—as an exploration of how information theory could be meaningfully applied to music composition (Hiller and Isaacson, 1959, p. 5). Though his methods paralleled that of others working in the nascent field of AI, his work was largely overshadowed by the more high-profile experiments of dedicated researcher within the cognitive sciences and engineering fields. [2] Hiller was more concerned with the importance of building a conceptual relationship between the content of the information and the method of generating musical scores. In later reflections on the nascent *Suite*, he claimed that he was disappointed by the acoustical analysis that lead, in his estimation, to fairly banal and conceptually empty results; the *Suite* began as an attempt to translate statistical analysis to *aesthetic* analysis by applying systems theory to musical historiography.

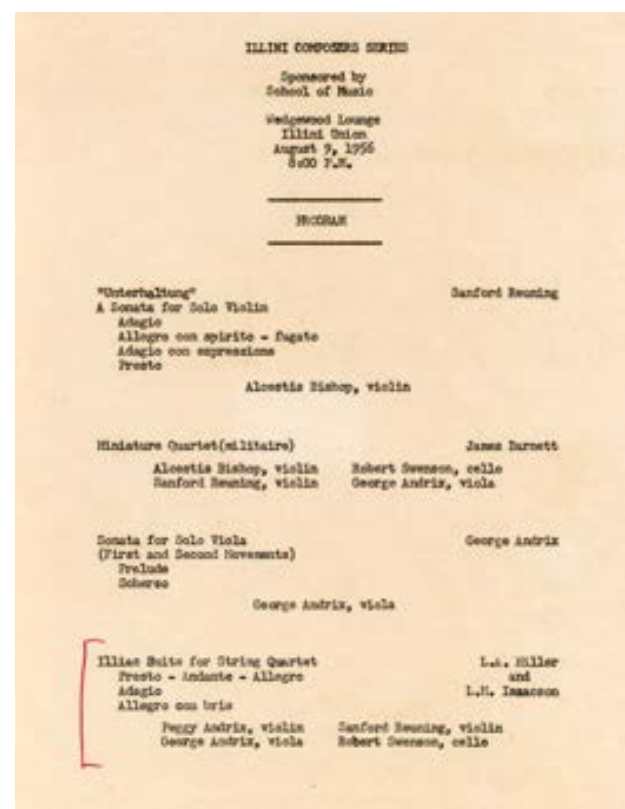


Figure 2: Original concert program for 9 Aug. 1956 performance of *Illiac Suite*. Image courtesy the University at Buffalo, SUNY Music Library.

Influenced by both the work of Warren Weaver as well as his time studying with composer Milton Babbitt at Princeton, Hiller focused on musical structural data to elevate the project from what he

considered the banal mimicry he observed in present probability table experiments (Hiller, 1959, p. 112). His definition of information theory, analyzing the oscillation between “randomness” (or “chance”) and “redundant” (“organized”) data within an entropic system, preceded this compositional process:

Information theory relates the “information content” of a sequence of symbols (be they letters of the alphabet or musical notes) to the number of possible choices among the symbols. Information content thus resembles entropy or the degree of disorder in a physical system. The most random sequence has the highest information content; the least random (or most redundant) has the lowest (Hiller, 1959, p. 109).

UPIA
(ILLIAC)
UP NEWS RELEASE
AUG. 10, 1954
CHAMPAIGN, ILL., AUG 10-(UPI)--A MUSICAL SUITE COMPOSED BY AN ELECTRONIC BRAIN WAS INTRODUCED BY A STRING QUARTET LAST NIGHT, BUT SOME LISTENERS DIDN'T DIG THE BEAT.
"WHY, IT DOES AWAY WITH THE NEED FOR HUMAN COMPOSERS," ONE WOMAN MUSIC LOVER REMARKED GLUMLY,
ANOTHER LISTENER SAID THE FINAL MOVEMENT OF THE COMPOSITION SOUNDED LIKE A "BARNYARD."
OTHERS AMONG THE AUDIENCE AT THE UNIVERSITY OF ILLINOIS APPEARED RESSENTFUL THROUGHOUT THE 15-MINUTE PERFORMANCE.
BUT PROF. ROBERT SWINSON, CELLIST AND LEADER OF THE QUARTET, SAID THE COMPOSITION, THE "ILLIAC SUITE" WAS "GOOD."
THE SUITE WAS COMPOSED BY ILLIAC, THE UNIVERSITY'S HIGH-SPEED DIGITAL COMPUTER.
THE COMPOSITION COVERED MUCH OF MUSICAL DEVELOPMENT FROM THE 16TH CENTURY ONWARD.
THE FIRST OF THE THREE MOVEMENTS WAS A SIMPLE PRESTO-ANDANTE-ALLEGRO WHICH COULD HAVE COME FROM THE PEN OF PALISTRINA. THE ADAGIO SECOND MOVEMENT LEANED TOWARD THE MODERN, AND THE THIRD ALLEGRO CON BRIO MOVEMENT WAS FRANKLY EXPERIMENTAL.
ILLIAC'S EFFORTS AT CLASSICAL MUSIC WERE SPONSORED BY L. A. HILLER, A CHEMIST-COMPOSER, AND L. M. ISAACSON, A RESEARCH ASSOCIATE.
HILLER TRANSLATED THE BASIC RULES OF COMPOSITION INTO ILLIAC'S LANGUAGE OF MATHEMATICAL SYMBOLS, THEN TAUGHT THE MACHINE TO PRODUCE RANDOM NOTES, REJECTING THOSE WHICH WOULD NOT FIT.
TP428A..

Figure 3: Copy of the United Press Release, reprinted in Bewley, J. (2004) “Lejaren A. Hiller: Computer Music Pioneer,” exhibition notes for Lejaren A. Hiller: Computer Music Pioneer, University at Buffalo, SUNY Music Library.

Rather than treating musical flourish as “noise” and rejecting this data as superfluous, Hiller’s stochastic approach preserved these details as behaviorally relevant and quantifiable. High information content allowing for randomness within the system—a capability only recently achieved through new advancements in probability theory and the use of a supercomputer—would not just allow for musical mimicry, but provide clues as to why aesthetic choices appear as they do throughout history:

The apparent paradox in this statement derives from the definition given the term ‘information’ in the theory. As Warren Weaver has observed, the term ‘relates not so much to what you do say as to what you could say’... The study of musical structures by information theory should open the way to a deeper understanding of the aesthetic basis of composition. We may be able to respond to Stravinsky’s injunction and cease ‘tormenting (the composer) with the why instead of seeking for itself the how and thus establish the reasons for his failure or success’ (Weaver 1949 cited in Hiller, 1959, pp. 109-110).

Many later descriptions summarize the *Illiac Suite* programming process as an attempt to simplify music composition by automating decisions conventionally made by the composer. For example, composer Gerald Strang’s analysis of the *Suite* in the 1969 *Cybernetic Serendipity* exhibition catalog suggests the process was solely an intermediary phase as Hiller explored options in automating music composition; he describes the *Illiac Suite* solely as a first step toward Hiller’s most well-known contribution to computer music, the assistive compositional software/programming language MUSICOMP (MUSIC Simulator-Interpreter for COMpositional Procedures) (Reichardt, 1969, p. 26).



Figure 4: The ILLIAC circa 1952. Image courtesy of the University of Illinois Archives.

However, Hiller’s *Illiac Suite* methodology was far more ambitious in its intent and conceptually rigorous in execution: by following Weaver’s insistence that information theory constituted the study of possibility, Hiller constructed probability tables derived from Renaissance and Classical models to contemporary serial compositions to aesthetically “train” each subsequent movement of the *Suite*. However, Hiller’s interpretive approach diverges most meaningfully from Mitchell’s contemporary ML probabilistic model in its

historical scope and intentionally limited scale; while contemporary machine learning applications often necessitate that researchers organize many, many more variables in probability distribution tables, Hiller's more interpretive, experimental scope—attempting to translate statistical analysis to aesthetic analysis—involved a much more intimate and carefully-constructed, historically-contingent set of parameters within the Western classical musical canon.

4. THE 4 MOVEMENTS OF THE ILLIAC SUITE

The *Suite's* movements corresponded to four experiments, each requiring new programs with increasingly complex screening rules exploring historical aesthetics. Hiller constructed the Suite to echo a historical progression from simple to complex melodies: the first mimicked Renaissance counterpoint rules, generating “simple” polyphonic melodies; the second produced four-voice segments within the confines of changing rules. Both the second and third experiments used a random chromatic method, expanding possible tonal values:

within the selected interval, no repetition or obvious patterns would be reproduced. The employment of the random chromatic method was meant to explore what Hiller identified as the major aesthetic difference between 17th and 20th century musical styles (Hiller and Isaacson, 1959, pp. 182-197; Hiller, 1959, p. 117).

For the third and fourth experiments, Hiller build his own stochastic process translating contemporary compositional rules to algorithmic systems that could represent the mounting complexity of serial compositions, particularly those explored by composer Arnold Schoenberg. Rather than rely on the more interpretive definition of “indeterminacy” espoused by his friend and future collaborator, John Cage, Hiller interpreted indeterminacy through adapting the MCMC methods he previously employed in chemical analysis research. This conceptual transformation of indeterminacy into a stochastic process—stringing randomly generated variables representing the present state to model how current changes affect future states—hinged on its ability to allow for *uncertainty*: the “stochastic” noise present in real-world examples, such as Hiller's analysis of the structure of music compositions throughout time (Hiller, 1959, pp. 23-24). This new set of randomly-generated variables disregarded any superfluous states, achieving “memorylessness.” Hiller and Isaacson generated integers sampled from a calculated probability distribution in the computer's memory until the machine saved a “melody,” designated “complete” after reaching a predetermined numeric length. The melody was printed on perforated tape, then hand-

transcribed into conventional musical notation (see Figure 5). In this way, the third and fourth experiments benefitted from MCMC memorylessness, mimicking contemporary compositions by eliminating probability favoring certain tones, automating Schoenberg's 12-tone technique preventing precedence of any of the 12 notes of the chromatic scale:

... the machine was first permitted to write entirely random chromatic music (including all sharps and flats)... With the minimal redundancy imposed by feeding in only four of the 14 screening instructions, the character of the composition changed drastically [from Experiment II]. While the wholly random sections resembled the more extreme efforts of avant garde modern composers, the later, more redundant portions recalled passages from, say, a [Bela] Bartok string quartet... The experiment concluded with some exploratory studies in Schoenberg's 12-tone technique and similar compositional devices (Hiller, 1959, p. 117; Schoenberg, 1975).




Figure 5: Printout from Lejaren Hiller Papers with edited *Illiatic Suite* manuscript. Images courtesy University at Buffalo, SUNY Music Library.

5. ASPIRATIONAL ARTIFICIAL INTELLIGENCE

Florian Cramer explains how computer code existed centuries before the invention of the computer in magic and musical composition, analyzing Pythagorean musicology through an arithmetic method of ascertaining the music of the spheres (Cramer, 2005). This kind of aspirational definition of mathematical musicology parallels machine learning researcher Zachary Lipton's description of AI as an impossibly lofty goal, “...

aspirational, a moving target based on those capabilities that humans possess but which machines do not” (Lipton, 2018). While this definition hints at divining a kind of psychological code rather than Pythagoras’ cosmological one, Tom Mitchell’s assertion that ML seeks to “build computer systems that automatically improve with experience” by ascertaining the laws that “govern all learning processes” constitutes a more pragmatic version of divining the secrets of the human psyche instead of the cosmos (Mitchell, 2016). Comparisons between such historically and culturally embedded practices and performances such as 17th century counterpoint compositions, Schoenbergian 12-tone experimental music, and the Bayesian methods should not assume a direct material relationship. However, the depth and breadth of the Hiller’s interpretive data sets and probability tables provides insight into the construction of predictive systems not only later computational compositions, but many of the ML algorithms for predictive modelling so commonly used today.

6. CONCLUSION

The *Suite* is far from the only experiment in stochastic, computational music exploration; composers and theorists such as Iannis Xenakis have produced stunning, forceful experiments in sound generation, spatialization, and orchestration. However, despite the nearly logarithmic growth in computing power and access to enormous data sets, Hiller’s attempts to ascertain stylistic meaning within these music composition training methods find no easy musicological research parallels. Hiller’s interest in systems theory and the interrogation of “style” or “aesthetics” still hangs loosely in the air, begging the question of whether stylistic analysis remains too general or banal a notion to concern serious musicological study.

In 2016, Sony CSL Research Laboratory produced the track “Daddy’s Car,” partially written by AI—the track’s harmonies and lyrics were composed by French musician Benoît Carré—in an attempt to recreate a Beatles-esque chart-topper. The track was produced using Flow Machines, an “augmented creativity” system that learns music styles from a huge database of songs by exploiting unique combinations of “style transfer, optimization and interaction techniques” (Sony CSL, 2016). Its GUI system uses patterned blob-like masses that shiver and pulse to differentiate between beat and lick “styles,” ready to be clicked-and-dragged for easy modular music composition. While Flow Machines may present a visually intriguing option for composing novices, its game-ified, pre-defined palettes mix metaphors of visual and sonic style, and greatly reduce the compositional results required by composers interested in the intricacies

and complexities of algorithmic music composition offered by older, non-ML-specific real-time audio synthesis programs such as SuperCollider (McCartney, 1996). The problem remains: can the flows of historical, musicological aesthetics/style be interrogated algorithmically in a way that is sensorially meaningful, and if so, is predictive analysis—employing the machine learning methodologies we so readily use for data analysis in every aspect of contemporary living—the way forward?

Hiller later reflected on the overall effect of the *Illiac Suite* in a *Scientific American* report, observing how the *Suite*’s “tonal” and “atonal” movements sounded similar despite their distinct algorithms; he acknowledged that the final movements were audibly indistinguishable from one another even though their computational processes greatly diverged. He proposed that its sound exceeded human perception, stating, “These correspondences suggest that if the structure of a composition exceeds a certain degree of complexity, it may overstep the perceptual capacities of the human ear and mind” (Hiller, 1959, p. 5). Though listening to the *Suite* proved challenging in its time, even infuriating to some, contemporary audiences of machine learning aesthetics have clearly developed an interest in mining the space between comprehension and complexity: glitch and dirty new media aesthetics, the fluid animation of Vuk Ćosić’s ASCII characters, and the remixing/modding of internet iconography by artists such as JODI (Joan Heemskerk and Dirk Paesmans) provide a wealth of artistic models of how complexity and human/computing limitations can be exploited to explore the making and deterioration of meaning through technology.

A particularly incisive visual example is MacArthur Genius Grant awardee Trevor Paglen’s installation “A Study of Invisible Images,” in which he uses computer vision to investigate “the collapsing distinctions between humans, machines and nature” (Metro Pictures, 2017). Contrasting this current (albeit well-placed) AI suspicion, curator Jasia Reichardt’s introduction to the 1968 *Cybernetic Serendipity* exhibition, in which the *Suite* was included, declares that the work “...deals with possibilities rather than achievements, and in this sense it is prematurely optimistic” (Reichardt, 1969, p. 5). Throughout his life and subsequent computer music experiments, Hiller remained more focused on the possibilities computer composition presented as a tool for historical analysis and pedagogy, in the fields of both musicology and information theory, evidenced by his founding of the electro-acoustic research facility Experimental Music Studios (EMS) at the University of Illinois in 1958, and his later research in the groundbreaking music composition programming language MUSICOMP (Hiller, 1969, pp. 71-83; Hiller and

Baker, 1964, pp. 62-90). Though he called the *Suite* “rather fragmentary” (Hiller, 1959, p. 5), a re-examination of Hiller's proto-machine learning compositional analysis methods reveals a rich historical vein for present machine learning acoustical research, presenting possibilities for contemporary scientists and musicians to critique aesthetic historical music stylistic conventions with the intent to adapt, and ultimately radically defy them.

7. REFERENCES

- Ariza, C. (2011) “Two Pioneering Projects from the Early History of Computer-Aided Algorithmic Composition,” *Computer Music Journal* 35 (3), pp. 40-56.
- Bewley, J. (2004) “Lejaren A. Hiller: Computer Music Pioneer,” exhibition notes for *Lejaren A. Hiller: Computer Music Pioneer*, University at Buffalo, SUNY Music Library.
- Cramer, F. (2005) *Words Made Flesh: Code, Culture, Imagination*. Rotterdam: Piet Zwart Institute.
- Dreifus, C. (1998) “A CONVERSATION: With Dr. Marvin Minsky; Why Isn't Artificial Intelligence More Like the Real Thing?” *The New York Times*, 28 July.
- Eckhardt, R. (1987) “Stan Ulam, John von Neumann, and the Monte Carlo Method,” *Los Alamos Science* (Special Issue), pp. 131-141.
- Encyclopedia of Statistical Sciences*, Volume 4, S. Kotz and N.L. Johnson (eds.) John Wiley and Sons, New York (1983)
- Hammersley, J. M., and Handscomb, D. C. (1964), *Monte Carlo Methods*. London: Chapman and Hall.
- Hiller, L.A. (1959) “Computer Music,” *Scientific American* 201(6), pp. 109-121.
- Hiller, L.A. (1969) ‘Some Compositional Techniques Involving the Use of Computers’ in von Foerster, H. and J. W. Beauchamp, *Music by Computers*. New York: John Wiley and Sons, pp. 71-83.
- Hiller, L.A. (1981) “Composing with Computers: A Progress Report,” *Computer Music Journal* 5(4), pp. 7-21.
- Hiller, L.A. and Baker, R. (1964) “Computer Cantata: A Study in Compositional Method,” *Perspectives of New Music* 3(1), pp. 62-90.
- Hiller, L.A. and Isaacson, L. (1959) *Experimental Music: Composition with an Electronic Computer*. New York: McGraw-Hill.
- Lipton, Z. ‘From AI to ML to AI: On Swirling Nomenclature & Slurried Thought,’ *Approximately Correct: Technical and Social Perspectives on Machine Learning blog* [online]. Available at: <http://approximatelycorrect.com/2018/06/05/ai-ml-ai-swirling-nomenclature-slurried-thought/> (Accessed: 5 June, 2018).
- Metro Pictures. (2017) “Trevor Paglen, A Study of Invisible Images, September 8 – October 21, 2017” press release. Available at: <https://www.metropictures.com/exhibitions/trevor-paglen4/press-release> (Accessed: 17 August, 2019).
- Mills, A. A. (1998) “Vermeer and the Camera Obscura: Some Practical Considerations,” *Leonardo* 31 (3), pp. 213-218.
- Mitchell, T. (1997) *Machine Learning*. New York: McGraw-Hill Science/Engineering/Math.
- Mitchell, T. (2006) “The discipline of machine learning” (CMU ML-06 108).
- Newell, A. and Simon, H. A. (1956) “The Logic Theory Machine: A Complex Information Processing System,” RAND Corporation report.
- Reichardt, J. (eds.) (1969) *Cybernetic Serendipity: The Computer and the Arts*. New York: Frederick A. Praeger.
- Robert, C. and Casella, G. (2011) “A Short History of Markov Chain Monte Carlo: Subjective Recollections from Incomplete Data,” *Statistical Science* 26 (1) pp. 102-115.
- Samuel, A. L. (1959) “Some Studies in Machine Learning Using the Game of Checkers,” *IBM Journal of Research and Development* 3 (3), pp. 210-229.
- Schoenberg, A. (1975) *Style and Idea*. Berkeley & Los Angeles: University of California Press.
- Turing, A. M. (1950) “Computing Machinery and Intelligence,” *Mind* 49, pp. 433-460.
- Winograd, T. (1971) “Procedures as a Representation for Data in a Computer Program for Understanding Natural Language,” *MIT AI Technical Report* 235.
- United Press (1956). (ILLIAC) CHAMPAIGN, ILL., AUG 10.

NOTES

1. Though Hiller and Isaacson's Illiac Suite is credited as the first computer-aided algorithmic composition, two additional experimental projects, developed independently in similar circumstances, came to be at approximately the same time: David Caplin and Dietrich Prinz's experiments with adapting Mozart's

Musikalisches Würfelspiel, and Sister Harriet Padberg's text-to-pitch mapping procedures (Hiller, 1970, p. 70; Christopher Ariza, 2011, p. 40). Hiller himself made note of Douglas Bolitho and Martin L Klein's Push Button Bertha (1956), as well as Richard C. Pinkerton's "tune maker," which involved constructing probability tables by averaging the probabilities using data from a collection of nursery tunes (Hiller, 1956, 112).

2. Foundational artificial intelligence research pioneered by Marvin Minsky and Terry Winograd at the Massachusetts Institute of Technology (MIT), Arthur Samuel at IBM, Allen

Newell, Herbert Simon, and Cliff Shaw at the RAND Corporation and Carnegie Mellon University (CMU) easily found its way into popular culture: Samuel's experiments with checkers-playing computers (Samuel, 1959), Newell, Simon, and Shaw's problem-solving Logic Theorist program (Newell and Simon, 1956), and Winograd's SHRDLU, a natural language understanding program (Winograd, 1971) garnered both scientific accolades as well as attention from the popular press. Minsky, who also published widely about the philosophy of AI, served as an advisor for Stanley Kubrick's 2001: A Space Odyssey (Dreifus, 1998).