Course: **DA3005 Independent Project (30 credits)**
2023
Master's Degree in Music, Electroacoustic Composition
Department of Composition, Conducting, and Music Theory

Text Supervisor: Annelie Nederberg

Muiredach O'Riain

The Well Trained Algorithm

An exploration of the use of AI as a tool for musical expression

Written reflection on independent work
This documentation also includes recordings and scores
of the following works: The Well Trained Algorithm

Kungl. Musikhögskolan

ABSTRACT

The Well Trained Algorithm is a piece that challenges prevailing conceptions of the use of AI tools in music through the reconceptualising of JukeBox, a generative AI model for music, as an instrument in its own right. Here, I am coining the term ‘instrumentisation’ to describe a methodology for applying the qualities and associations of a musical instrument to a traditionally non-musical object. To showcase this conceptual approach, this model of thinking is applied to aid in the composition of the AI-generated musical piece, *The Well Trained Algorithm*. Through this process of ‘instrumentisation’, musical terms such as tuning and timbre are redefined to better relate to the specific affordances of an artificially intelligent system. The composition is informed, then, by an exploration of a system's instrumental possibilities, leading to a more effective and artistic use of the technology in the creative process. The seminal works of J. S. Bach and La Monte Young, *The Well Tempered Clavier* and *The Well Tuned Piano*, respectively, provide a historical, musical, and theoretical context for the piece as well as the datasets used to fine-tune the JukeBox model. With this thesis, I ask if, through a process of ‘instrumentisation’ AI technology can be successfully reconceptualised as a musical instrument as a means to promote artistic expression.

Keywords:
Artificial Intelligence (AI), Machine Learning (ML), Generative Music, J. S. Bach, LaMonte Young, *The Well Tempered Clavier*, *The Well Tuned Piano,* Tuning System, Instrumentisation

Table of Contents

INTRODUCTION

The Well Trained Algorithm is a project that aims to challenge prevailing notions of artificial intelligence (AI) in music by reconceptualising an AI system as an instrument in its own right through a process of 'instrumentisation'. This thesis will begin with a theoretical contextualisation of the use of AI as a creative tool, exploring its current perception in society and its use in the creation of this piece. To provide a conceptual framework for the proposed approach of "instrumentisation," the author will then offer a theoretical context in which an abstract object can be redefined in relation to its use as a musical instrument. By identifying the current modes of thought surrounding AI technology, the thesis will propose the concept of 'instrumentisation' as a means to break free from these models of thinking to better inform the use of a generative AI as an artistic tool to aid in a creative practice. To demonstrate the application of this approach, this thesis will present a musical composition, *The Well Trained Algorithm*.

Artificial Intelligence

"Artificial intelligence is not, by definition, simulation of human intelligence" ~ John McCarthy[1]

According to the definitions proposed by Russel and Norvig, artificial intelligence (AI) describes the ability of a machine to act as an intelligent agent. That is to say, to act autonomously, to consider its environment and to make rational, goal-oriented decisions. In the field of AI, intelligence is defined as the autonomy of a system to enact goal-directed behaviour towards the fulfilment of an objective function. In addition, they state that an intelligent agent may also learn to perform a function more efficiently through a process of knowledge acquisition known as machine learning (ML).[2]

AI itself is not a new concept. As early as 1947, British mathematician and one of the founding fathers of modern computing, Alan Turing, is said to have described" a machine that can learn from experience" and a year later, in 1948, he expanded upon this idea in a report entitled *Intelligent Machinery* which outlined many of the central concepts of AI relevant to this

[1] J. McCarthy, AI@50 Conference, *What Was Expected, What We Did, and AI Today*, Filene Auditorium, Moore Hall, Dartmouth College, Hanover, NH, July 13, 2006, cited in conference notes by M. H. Maker (2006) *AI@50: AI Past, Present, Future.* Archived at https://web.archive.org/web/20081008120238/http://www.engagingexperience.com/2006/07/ai50_ai_past_pr.html on archive.org, (archived from the original on January 3, 2006), (retrieved October 16, 2008), (accessed 5 June 2023).
[2] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Pearson Series in Artificial Intelligence, 4th US edn., Hoboken, Pearson Education Limited, 2022, pp. 1-62.

day.[3] During the 1950s, ML would be developed through applications in the field of game theory, with board games such as chess and checkers providing a robust ruleset and objective function. In 1955, Arthur Samuel programmed an IBM 704 computer to run the first ever checkers-playing algorithm that utilised a loss function to teach itself how to play, marking the first implementation of ML algorithms for complex problem solving. However, the question remains: how intelligent are these machines actually?

Turing would follow up his earlier work with *Computing Machinery and Intelligence,* wherein he begins by asking the question, "Can machines think?"[4]. To address this hypothesis, Turing proposed *The Imitation Game*, a test to assess whether a machine could demonstrate intelligent behaviour that is comparable to, or indistinguishable from, that of a human. Turing suggested that a human evaluator judge text-based conversations between a human and a machine, where both are hidden from the evaluator. If the evaluator was unable to distinguish between the two, the machine was deemed to have passed the test, regardless of the accuracy of its answers.[5] However, there has been some debate over the legitimacy of this test regarding whether or not a machine can really be described as thinking.

In the case of a "weak AI", as proposed by philosopher John Searle, the machine is not actually aware of the conversation it is having; its responses are informed by a predetermined chain of probabilities and not by any deeper 'intelligent' understanding.[6] In this case, the AI passes the test, but it does so without thinking, at least not in the way that we commonly understand. Searle states that here the machine is instead demonstrating its ability to convincingly simulate the process of thinking for a human observer.[7] Taking a Cartesian approach to Searle's analysis of this problem of computer thought, we may assume that, as individuals, we understand our own existence as that of a thinking entity; *cogito, ergo sum*. However, when conversing with another entity, human or otherwise, though we may assume a separate intelligence we cannot know for certain. As Turing put it, "it is usual to have a polite convention that everyone thinks"[8]. Here, Turing suggests that what actually determines the efficacy of the AI is whether or not the human observer is convinced by its responses, it is the context of our interaction with the machine that provides its usefulness. In this thesis, I approach the use of AI in this way, treating the AI not as a separate creative intelligence per se, but rather as a creative tool, and define a system of thinking in order to better work with this technology and create a convincing musical output: *The Well Trained Algorithm*.

---

[3] B. J. Copeland, "artificial intelligence". *Encyclopedia Britannica*, 1998, https://www.britannica.com/technology/artificial-intelligence, (accessed 5 June 2023).

[4] A. M. Turing, 'I.—COMPUTING MACHINERY AND INTELLIGENCE', *Mind*, vol. 59, no. 236, 1950, pp. 433–460, https://academic.oup.com/mind/article/LIX/236/433/986238, (accessed 5 June 2023).

[5] A. M. Turing, 'I.—COMPUTING MACHINERY AND INTELLIGENCE', *Mind*, vol. 59, no. 236, 1950, pp. 433–460 https://academic.oup.com/mind/article/LIX/236/433/986238, (accessed 5 June 2023).

[6] J. Searle, "Minds, brains, and programs." *Behavioral and Brain Sciences*, vol. 3 no. 3, 1980, pp. 417-424.

[7] K. Arkoudas and S. Bringsjord, K. 'Philosophical Foundations' in Frankish and W. Ramsey (Eds.), *The Cambridge Handbook of Artificial Intelligence,* Cambridge University Press, 2014, pp. 34-63.

[8] A. M. Turing, 'I.—COMPUTING MACHINERY AND INTELLIGENCE', *Mind*, vol. 59, no. 236, 1950, pp. 433–460 https://academic.oup.com/mind/article/LIX/236/433/986238, (accessed 5 June 2023).

In the 70 years since the publication of Intelligent Machinery, AI has been steadily expanded for use in countless fields, including medical diagnoses, image recognition, stock trading, language translation, and robot control, to name a few. Recently, during the COVID-19 pandemic, AI was extensively utilised by researchers to "accelerate COVID-related drug discovery”[9]. In the past decade, however, an “AI spring”[10] has ushered in an unprecedented period of AI development. From self-writing emails to self-driving cars, artificially intelligent technology has become increasingly present in our everyday lives. Historically, this technology has been used primarily in an industrial context, often as a predictive tool for logistics, targeted marketing campaigns, and e-commerce. However, by providing different input data for these machine learning models, it is possible to train an AI to generate what can be considered an artistic or creative output in the form of text, images, or audio.

In the field of art specifically, these technologies have become increasingly sophisticated in their ability to emulate certain styles and produce convincing approximations of human artistic expression. Projects such as MidjourneyV5 and Stable-Diffusion are now able to create hyper-realistic AI-generated "photos" from word prompts, and the widespread availability of these tools has led to several viral deep-fake hoaxes, such as hallucinogen-inspired depictions of Pope Francis wearing a big coat.[11] The increased media attention in the aftermath of each hoax has led to many discussions on the definitions of art or the ethical use of AI, and the debate over whether computer-generated works can even be considered art in the first place is still ongoing. Within this discourse, I have identified two main perceptions surrounding the use of AI in art that I aim to challenge with this project.

One concern about generative AI is its perceived inability to offer original insights or expression compared to human-driven methods. Critics such as American linguist Noam Chomsky have argued that in this reliance on statistical techniques, AI technology lacks any form of "critical experiment"[12] and therefore does not meet the standards of a scientific discipline. In this way, Chomsky maintains that AI is unable to offer any explanatory insight or original knowledge.[13] Indeed, at its core a generative AI identifies patterns within existing datasets, which can result in a tendency to recreate existing patterns rather than offer novel outputs. In discussing the use of AI in art, this argument is often used to dismiss the output of

[9] D. Zhang et al., *The AI Index 2021 Annual Report,* AI Index Steering Committee, Human-Centered AI Institute, CA, Stanford University, Stanford, 2021, p. 2., https://aiindex.org/2021-report/, (accessed 5 June 2023).
[10] J. Manyika and J.Bughin, ‘The coming of AI Spring’, *McKinsey Global Institute,* Project Syndicate, 2019, https://www.mckinsey.com/mgi/overview/in-the-news/the-coming-of-ai-spring, (accessed 5 June 2023).
[11] D. Bennett, ‘AI Deep Fake of the Pope’s Puffy Coat Shows the Power of the Human Mind’, *Bloomberg,* Bloomberg, 2023, https://www.bloomberg.com/news/newsletters/2023-04-06/pope-francis-white-puffer-coat-ai-image-sparks-deep-fake-concerns, (accessed 5 June 2023).
[12] ‘#53 Noam Chomsky: Language, Cognition, and Deep Learning’ , *Lex Fridman Podcast,*[podcast/online video], Presenter L. Fridman, interview with Noam Chomsky, U.S, Lex Fridman, 2019, https://www.youtube.com/watch?v=cMscNuSUy0I, (accessed 5 June 2023).
[13] Y. Katz, ‘Noam Chomsky on Where Artificial Intelligence Went Wrong’, *The Atlantic*, The Atlantic, 2012 https://www.theatlantic.com/technology/archive/2012/11/noam-chomsky-on-where-artificial-intelligence-went-wrong/261637/, (accessed 5 June 2023).

these models as overly derivative.

Secondly, there are concerns surrounding the perceived threat of AI as a displacement of the human. As fears of job loss via automation sow a form of AI anxiety among the public, that the use of AI will lead to a "dystopian future"[14]. This may be partly influenced by the portrayals of AI in popular fiction, where corrupted or malevolent AI often plays an antagonistic role, acting in direct opposition to humanity. An analysis of the depiction of AI in films such as 2001: A Space Odyssey, The Terminator, and The Matrix reveals a common thread that can also be seen in current perceptions of AI: the threat of displacement by a technologically advanced other. Recently, this fear was realised when an AI-generated image won first prize for digital arts at the Colorado State Fair in September 2022, leading many to criticise the artist's use of AI, deeming it deceptive and suggesting it was unfair to "beat out other human artists with something he asked a machine to create".[15]

These stories contribute to the sensationalisation of AI in the media, where the technology is often understood through the lens of its technological limitations when compared to human intelligence. With this perception, a false dichotomy is created whereby artificial creativity is viewed as an alternative to human creativity. In an attempt to evaluate the relative merits of AI within a purely oppositional framework, we limit our understanding of its potential as a creative tool within a greater artistic process. Through an analysis of the use of AI in the composition of *The Well Trained Algorithm*, this thesis aims to challenge these notions of AI as either a derivation of or alternative to human expression by redefining the relationship between the artist and this technology. I propose that through a process of 'instrumentisation' it becomes possible to identify the specific artistic properties of the AI system and to work with this technology instead as an artistic tool to aid in the creation of a new and interesting musical expression.

Instrumentistion

'Instrumentiastion' is a term that I will use throughout this thesis, both as a concept and a proposed method of working more effectively with this emerging technology. Here I present a definition of this term and how it relates to the use of AI in the composition of *The Well Trained Algorithm*.

As a starting point, I borrow the concept of affordances as proposed by American psychologist James J. Gibson to begin to offer an alternative conceptualisation of AI as an artistic tool. In *The Ecological Approach to Visual Perception*, Gibson introduces the idea of a

---

[14] M. Iyanda, 'ChatGPT: How the World Is Reenacting Moral Panics on Artificial Intelligence', *Tekedia,* Tekedia, 2023, https://www.tekedia.com/chatgpt-how-the-world-is-reenacting-moral-panics-on-artificial-intelligence/, (accessed 7 June 2023).
[15] R. Lerman. "At the Colorado State Fair, visitors get a glimpse of A.I.'s mid-journey.", *The Washington Post*, The Washington Post, 2022, https://www.washingtonpost.com/technology/2022/09/02/midjourney-artificial-intelligence-state-fair-colorado/, (accessed 5 June 2023)

"value-rich ecological object" (VREO), where the benefits and drawbacks of an object are considered in relation to the organism that interacts with it. This approach goes beyond the value-neutral language of physics and highlights how an affordance is well-suited or ill-suited to a user's needs.[16] In essence, an affordance reflects the beneficial or harmful aspect of objects, in relation to its use. I propose that by examining the affordances of AI in relation to its use as a musical object, we are able to utilise this technology more effectively as a tool for artistic expression.

In order to redefine a generative AI as a VREO, I first objectify the AI by removing the abstraction of the algorithm and considering it as a physical entity. Next, I define the user as a composer, and the intention of that user as the creation of musical sound. From here, I can begin to assign affordances to the AI object to better understand its value in relation to this use. However, through 'instrumentisation' I propose that this line of thinking can be taken one step further. By considering the AI not as an algorithm or a VREO in the abstract but as a musical instrument in its own right. Through the reinvention of the AI as a musical instrument, we can now apply musical and instrumental concepts such as pitch, tone, tuning, and timbre as a way to more intuitively describe and understand the AI as creative tool.[17] It is this process of redefining the AI as a musical tool specifically and thereby placing it within the framework of a greater musical context that I will refer to as 'instrumentisation'.

Another aspect of 'instrumentisation' is that it helps to create a relational model of value between the AI instrument and other musical instruments within a greater ecosystem. Here 'instrumentisation' acts as the obfuscation of the machine as proposed by the Turing test, levelling the field by removing the label of AI from the instrument. In doing so it becomes easier to form an objective comparison between the specific instrumental affordances of the AI instrument and the affordances of other musical instruments, as a way to better understand its value as a creative object for a specific musical purpose.

From a creative standpoint, we can also utilise this model to identify the musical affordances of an AI instrument at a higher level and, from there, begin to explore and redefine these limits through our interaction with the instrument. In this way, I believe that 'instrumentisation' can inform a more original and creative use of this technology in a musical context and, in doing so, challenge the aforementioned perceptions that surround the use of AI in music. Over the course of this thesis, I will address the following research question as a means of analysing the efficacy of this process: Through 'instrumentisation' can AI technology be successfully reconceptualised as a musical instrument as a means to promote artistic expression in the composition of a musical work?

[16] J. Gibson, *The Ecological Approach to Visual Perception*, New York. NY, Psychology Press, Taylor & Francis Group, LLC, 1986, pp. 127-146.

[17] I. Héroux and M. Fortier, *Expérimentation d'une Nouvelle Méthodologie Pour Expliciter Le Processus de Création d'une Interprétation Musicale*, Les Cahiers de La Société Québécoise de Recherche En Musique, 2015, pp. 67–79.

METHOD

To showcase 'instrumentisation' as a means of informing the creative process *The Well-Trained Algorithm* takes this model of thinking and applies it to the conceptualisation and composition of a piece of AI-generated music. To achieve this, I have chosen to work with JukeBox, a generative model for music that combines AI and machine learning techniques to synthesise audio. To strengthen this conceptualisation I use the music of Bach and La Monte Young, not only to provide a historical and musical context for the work but also as discrete datasets to train the AI model. To begin, I will provide a summary and analysis of the most important technology and musical works that were utilised in the composition of *The Well Trained Algorithm.*

Technology

JukeBox

JukeBox is a generative machine learning model developed by OpenAI. It was released to the public in November 2020 as the result of training a deep learning algorithm on a dataset of 1.2 million songs. When prompted, JukeBox is able to generate high-quality audio in the style of a wide range of musical artists and genres.[18] Essentially, it works by analysing a dataset of music and, over time, defining an algorithm, or a set of instructions, that a computer can then use to automatically synthesise new music in a similar style.

Fine Tuning

The initial OpenAI JukeBox model can be manually trained further to get much more coherent results through a process called fine tuning. Fine tuning the model on a smaller data set of specifically selected music allows for a greater level of control over the final output audio.

Google Colab + Google Drive

Google Colab is a Google Research product designed for machine learning and data analysis. It's a Jupyter notebook service that runs Python code in a browser and can access and edit data through Google Drive. Runtimes can be allocated resources, such as GPUs, to increase computational power. In this project, Colab was used as a platform to both fine tune the JukeBox algorithm and prompt the model to generate audio.

[18] P. Dhariwal, et al., 'Jukebox', *OpenAI*, San Francisco. CA, 2020, https://openai.com/research/jukebox, (accessed 5 June 2023).

Training the AI

In order to reject the prevailing notions of artificial intelligence and reconceptualise JukeBox, we must not only utilise the AI through the lens of 'instrumentisation', but also communicate the concept of 'instrumentisation' as an essential part of the composition process. Since *The Well Trained Algorithm* is presented as a musical piece, the question becomes: how can this process be expressed within the context of a musical composition?

Extra-Musical characteristics of Sound

In a musical context without lyrics, it has been argued that intention is usually communicated through one of two primary methods: explicitly through the perceptual attributes of the objective musical material, as put forth in Austrian philosopher Eduard Hanslick's foundational analysis of the subjective aesthetics of music, *Vom Musikalisch-Schönen*[19], or implicitly through contextual associations and the reference of extra-musical components, as defined in Leonard B. Meyer's *Emotion and Meaning in Music*.[20] In a musical context, without lyrics, it has been argued that intention is usually communicated through one of two primary methods: Explicitly through the perceptual attributes of the objective musical material as put forth Austrian philosopher Eduard Hanslicks foundational analysis of the subjective aesthetics of music. Explicit features are perceptible, quantifiable; they include things like volume, melody, pitch, and timbre. Implicit extra-musical characteristics are subjective qualities of sound that indirectly relate to the music. These immeasurable properties reflect the personal priorities of the musician or the piece's situation within a musical context. These salient aspects of music are often described in terms of style, expression, nuance, and emotion.[21] Having identified these modes of communicating intention, we apply our conceptual model to analyse the specific musical affordances allowed by the JukeBox algorithm and evaluate how to effectively utilise these affordances to communicate the idea of 'instrumentisation'.

In the preface to *Musical Instruments: history, technology, and performance of instruments of western music*, it is put forward by the authors Campbell, Greated, and Myers that "The primary function of all musical instruments is to generate musical sound".[22] This holds true for JukeBox; however, while traditional instruments allow performers to control the objective musical features of their output, JukeBox does not. Once a model is trained, the audio output is generated algorithmically. As a black-box system, its musical process can only be viewed in terms of its inputs and outputs. To convey meaning, then, JukeBox must rely solely on the implicit aspects of music that can be imparted from its training data.

[19] E. Hanslick., *On the musically beautiful*. Translated by G. Payzant. Cambridge. MA, Hackett Publishing, 1986.
[20] L. B. Meyer, *Emotion and Meaning in Music,* Chicago: Chicago University Press, 1956.
[21] I. Héroux, 'Creative processes in the shaping of a musical interpretation: a study of nine professional musicians' *Front Psychol*, vol. 9, no. 665, 2018, https://www.frontiersin.org/articles/10.3389/fpsyg.2018.00665, (accessed 5 June 2023), p. 4.
[22] M. Campbell, C. A. Greated, and A. Myers, *Musical Instruments: history, technology, and performance of instruments of western music*, Oxford University Press, 2004, p. v.

It is here that JukeBox affords us an advantage over other generative AI models that use a symbolic representation and a secondary sound source to generate music. Though MIDI files may allow a composer more control over the objective characteristics of their sound through the use of synthesis, JukeBox's generative algorithm analyses the raw audio of a recording, bypassing the symbolification of the input data and allowing it to capture and reproduce the extra-musical qualities of a performance that might otherwise be lost.[23] In selecting appropriate input data, we can utilise this aspect of the JukeBox algorithm to help to create musical parallels between *The Well Trained Algorithm* and its training data by allowing these extra-musical aspects to place its output within a predefined musical identity. Through incorporating works in the training data that deal with the concepts of instrumentality and affordances addressed in this thesis, the extra-musical characteristics of these works imparted to the output of the JukeBox algorithm can help to provide a musical and theoretical context for the piece and establish a musical framework that can be utilised to reject the current perceptions of AI.

Choosing the Training Data

In order to work with JukeBox and communicate these concepts effectively, I chose to begin by training the JukeBox algorithm on two specific works: J. S. Bach's *The Well Tempered Clavier* (TWTC) and La Monte Young's *The Well Tuned Piano* (TWTP). In naming the piece *The Well Trained Algorithm*, I am making a conscious reference to these two specific musical works while also contextualising *The Well Trained Algorithm* as a continuation of the ideas explored by Bach and Young (Young's naming of TWTP is itself a reference to Bach's TWTC). By training the JukeBox algorithm on these pieces specifically, the intention is to draw from this lineage of works in order to communicate the aim of this project: To redefine our understanding of artificial intelligence within an instrumental framework, in order to better understand its musical affordances, specifically through 'instrumentisation', TWTA relates the fine tuning of the AI model to the tuning of a traditional instrument, and devises a radical approach to tuning as a means of examining the overlooked potential of an instrument. Here I provide a summary of these works, followed by an explanation of their role in this project.

J. S. Bach - The Well Tempered Clavier

*The Well Tempered Clavier (*TWTC), or *Das Wohltemperierte Klavier* in the original German, is a collection of two sets of 24 preludes and fugues written for clavier (any instrument with a keyboard) by Johann Sebastian Bach. These pieces are collected in two books, with each set containing a prelude and fugue in every commonly used major and minor key, starting with Prelude No. 1 in C major and working up through the octave, finishing with Fugue No. 24 in B minor. It is believed that Bach wrote TWTC to demonstrate the advantages of an alternative "well-tempered" tuning system that allowed for music to be played in any major or minor key without sounding perceptibly out of tune.[24]

[23] P. Dhariwal, et al., 'Jukebox', *OpenAI*, San Francisco. CA, 2020, https://openai.com/research/jukebox, (accessed 5 June, 2023).
[24] J.S. Bach and W.A. Palmer, *The Well-Tempered Clavier, Volume I*, Van Nuys. CA, Alfred Music, 2005, p. 5.

La Monte Young - The Well Tuned Piano

The Well Tuned Piano (TWTP) is an ongoing musical work by American composer La Monte Thornton Young. The piece is played on an Imperial Bösendorfer Grand piano specifically intoned to a tuning system developed by Young. Initially conceived in 1964, the tuning is centrally important to the piece and was developed by working with principles of just intonation, where the intervals between notes are based on simple integer ratios. Debuting live for the first time ten years after its initial conception in 1974 at around 45 minutes long, The Well Tuned Piano would evolve over the coming decades into a longform semi-improvisatory solo piano performance, which can take upwards of 5–6 hours to perform.[25][26]

Training Data and Fair Use

The recordings of TWTC utilised in the creation of this piece were taken from recordings of Bach's 48 by Canadian pianist Glenn Gould, recorded on a Steinway CD 318 grand piano between 1964 and 1971, and the recording of TWTP is from the Gramavision release of *The Well-Tuned Piano 81 X 25 6:17:50 - 11:18:59 PM NYC* (US, 1987) as played by Young.

At the time of writing this thesis, the laws surrounding the use of copyrighted music to train an AI lie within a grey area of legislation. Since the works mentioned here were not directly sampled in any way in the creation of TWTA, and due to the transformative nature of the work on top of this being a non-commercial project, there is a strong case to make for the use of these pieces in this project falling under fair use. However, whether or not using music in this way to train AIs will continue to be protected going forward is yet to be seen. Currently, though there is no legal precedent, "Universal Music Group (UMG), who controls approximately a third of the global music market, has asked streaming services to obstruct AI companies from creating "new" music using copyrighted music".[27] Furthermore, it is unclear to whom the rights of TWTA would or should actually belong. As a legislature develops to deal with these emerging technologies and their consequences, I feel that it is important to establish a language around the role of AI as a creative tool. Drawing a parallel between AIs and instruments could be one way to better define and deal with the problems of ownership that currently surround AI-generated music.

---

[25] 'Interview with La Monte Young and Marian Zazeela', *KPFA-FM* [radio broadcast], Presenter R. Jennings, Berkeley. CA, KPFA, 1984, https://archive.org/details/AM_1984_08_XX, (accessed 5 June 2023).
[26] T. Riley, D. Wolf, L.M. Young, M. Zazeela and D. Farneth, *The Well-Tuned Piano 81 X 25 6:17:50 - 11:18:59 PM NYC* [LP liner notes], Gramavision Records, New York. NY, 1987, p. 9.
[27] J. Aswad, 'Universal Music Asks Streaming Services to Block AI Access to Its Songs', *Variety*, Variety, 2023, https://variety.com/2023/music/news/universal-music-streaming-services-block-ai-1235582612/, (accessed 25 May 2023).

Identifying Affordances and Establishing Parallels

Both TWTC and TWTP are ground-breaking musical works in their own right, but more explicitly, they are works that expose and explore the limits imposed by the specific affordances of their instruments, the clavier and the piano. The keyboard layout of a clavier affords a player the ability to simultaneously sound multiple notes with each hand. Bach takes full advantage of this affordance with his contrapuntal music, as the complicated interplay of melodic lines that define this style of composition would be near impossible to accurately recreate on a single-voiced instrument. However, a keyboard interface also has its drawbacks.

By the nature of its design, one key cannot be assigned to more than one note, and once that note is tuned, it cannot be easily re-tuned during play. Here, unlike the fingerboard of a violin or the continuous handslide mechanism of a trombone, where the pitches allowed to a player are bound only by the physical limits of the instrument and the player, the clavier restricts the tones that a player is able to produce to twelve per octave. This restriction can also be understood as an affordance.

In order to facilitate harmony while taking advantage of the polyphony afforded by a clavier, each key is tuned to a note of a specific frequency, where the interval between each note is relative to every other note. However, if we consider the tuning of a modern piano to be solely determined by the relationship between these frequencies, we are only seeing one side of the story. In practise, the tuning of each note on a piano is decided not only by their relative frequencies but also by the specific physical affordances of the piano as an instrument, which are in turn influenced by the pre-existing musical technologies as well as traditions of harmony and intonation. For example, why is it that the octave on a piano is split into twelve notes as opposed to, say, the seventeen notes per octave seen on a traditional Turkish Bağlama or vice versa? Each instrument can also be understood in the context of its creation. Regarding the layout of a clavier, the 12-key octave can be traced back to the development of chromaticism and the twelve-tone system in western music over one hundred and fifty years prior to the invention of the first piano. Over time, these different relationships develop to form a musical ecosystem that influences the way music is created and experienced. With TWTC and TWTP, Bach and Young are able to conceptualise tuning both as an inherent property of their instrument and within the context of a greater musical tradition. With this understanding, they are able to then explore the limits of these affordances and redefine them to break from tradition and offer an alternative.

Training the Algorithm

"You pushed the button and out came hundreds and thousands of sonatas or whatever"~ David Cope [28]

For both the TWTC and TWTP, I argue that the initiation of the creative process was in identifying a problem—a limit imposed by the current system of musical expression. For Bach, the dissonances arising from the prevailing meantone temperament imposed a limit on the range of keys offered when playing on a clavier. In identifying this limit, Bach could then explore a way to bypass it within the affordances offered by the instrument. Well-tempering solves this issue by implementing a more relative system of harmony, allowing for each key to be utilised.

However, tempering compromises any pure or just harmony between individual keys in favour of a relative global tuning. Young felt that this method of intonation didn't afford the ability to create and compose music with the pure harmonic intervals that he believed were the key to an alternative form of musical expression. In the case of TWTP, Young goes a step further than Bach, devising his own method to derive and so describe a new series of pure frequency relations with which to express himself, a new tuning system.

If we consider the JukeBox AI through the framework of 'instrumentisation' we can reconceptualise the training data itself as a tuning system, as it is through the analysis and recreation of patterns identified in this data that the model learns to generate music. During the training process, the weights and biases that inform the output of the algorithm are adjusted to better generalise the information in the training data set. However, this generalisation of data can be problematic, and herein lies the limit of the JukeBox tuning system as imposed by the prevailing system of temperament today.

Due to the almost universal use of twelve-tone equal temperament (12TET) in western popular music, the vast majority of music that the JukeBox algorithm is initially trained on is recordings of instruments tuned to a 12TET standard.[29] The result is an algorithm with a learned output that recreates the same frequency relationships of this music. Essentially, the JukeBox algorithm, despite not having any physical tuning apparatus, is tuned in 12TET. With the Well-Trained Algorithm, I propose that this limit is imposed not by the affordances of the AI, which in theory should be able to recreate musical notes in any tuning system. Rather, this is a limit imposed by the prevalence of these systems in modern popular music and can be combated by our approach to interacting with the AI system.

[28] D. Cope, David Cope: 'You pushed the button and out came hundreds and thousands of sonatas' [interviewed by T. Adams], 11 July 2010, *The Guardian*, The Observer, 2010, https://www.theguardian.com/technology/2010/jul/11/david-cope-computer-composer, (accessed 6 June 2023).

[29] P. Dhariwal et al., 'Jukebox: A Generative Model for Music', *ArXiv*, vol. abs/2005.00341, 2020, p. 11. http://arxiv.org/abs/2005.00341, (accessed 20 April 2023).

In order to break from this tradition and offer an alternative, as done by Bach and Young in their respective works, we too must implement a new tuning system. Here, the solution is to fine-tune the JukeBox model by continuing to train it on new music that doesn't fit within this system. Through this fine tuning, we essentially create a new instance of the JukeBox algorithm with a different tuning system capable of producing different pitches and timbres to those of the original algorithm.

The Tuning Process

In deviating from an established standard, we are forced to begin with the standard and then alter it in some way. In order to begin the composition real, the proposed system must be actualised through a process of experimentation and re-calibration, a realisation of one's departure from the biases imposed by a predefined tradition.

This means that we must first run a fine-tuning program to change the tuning of JukeBox based on the music of Bach and Young. Theoretically, the more that the algorithm is trained, the closer it will get to simulating the data set exactly. This process can be thought of much like the "perfecting"[30] of the tuning over weeks prior to a performance of "*The Well Tuned Piano*". Young describes, the goal of tuning to render these pitches as accurately as possible "inasmuch as in this case the intervals for The Well-Tuned Piano must actually be tuned within the limitations imposed by the in-harmonicities inherent to stringed instruments."[31] Where the retuning of *The Well Tuned Piano* has a pitch as a goal, JukeBox uses digital audio to train itself. However, without first collecting and processing this training data into a usable corpus, the model cannot be finetuned, as there is no goal for it to orient itself towards. The final preliminary step before tuning, is to source and process the training data that will become our new tuning system, in order for the tuning process to be effective.

Gathering and Processing the Data

As the JukeBox training algorithm can only process audio files up to a certain length, I moved each recording to a folder on a hard drive and processed the contents each folder with ffmpeg, a command line media converter tool. A write a command to split any longer files into shorter 5 minute sections, rename each separate part and convert the files to a .wav format. These segmented versions of the recordings were uploaded to a google drive folder to be analysed by the JukeBox algorithm.

---

[30] T. Riley, D. Wolf, L.M. Young, M. Zazeela and D. Farneth, *The Well-Tuned Piano 81 X 25 6:17:50 - 11:18:59 PM NYC* [LP liner notes], Gramavision Records, New York. NY, 1987, p. 8.
[31] T. Riley, D. Wolf, L.M. Young, M. Zazeela and D. Farneth, *The Well-Tuned Piano 81 X 25 6:17:50 - 11:18:59 PM NYC* [LP liner notes], Gramavision Records, New York. NY, 1987, p. 5n1.

Tuning/Training

To facilitate the training of the JukeBox algorithm I utilised Google Colab to run an edited version of the "Zags_1bl_finetuning_l2_labeled_with_lyrics____ft5"[32] a python notebook, uploaded by user "Zags" on the JukeBox Community Discord (an instant messaging board). This notebook downloads the JukeBox algorithm and allows it to finetune on audio files stored in a folder on Google Drive. The new, fine-tuned algorithm is then stored as a ~5GB .tar file in another folder on Google Drive, which is continually updated as the algorithm is trained farther on the new tuning system.

Sampling/Playing

Once the new instance of the JukeBox algorithm has been sufficiently retuned the file containing this re-tuned version can then be accessed with another notebook in order to generate audio. For this step, which I refer to as "sampling" I used an edited version of Zags "Zags_5_0_ft4_multi17_release" notebook, to interact with the newly tuned algorithms.

Composition

With TWTC and TWTP where the composers utilised traditional instruments, a full re-tuning of their instrument was carried out before the piece was played. Due to the unique affordances of working with an AI instrument, JukeBox can be both "tuned" and "untuned" simultaneously and at certain points in the composing process, "re-tuning" and playing happen concurrently. As such this process of "playing" the AI cannot be described in a traditional sense, so here I provide a written score, to outline the means of recreating *The Well Trained Algorithm* with the JukeBox algorithm.

[32] Zags_1bl_finetuning_l2_labeled_with_lyrics___ft5.ipynb' [source code], Zag, 2022, https://colab.research.google.com/drive/17Uk7R2nG2I5W-unOX2uU576wGdfvI_13, (accessed 6 June 2023).

SCORE

Preface

The Well Trained Algorithm is a piece of music made for and played entirely by the JukeBox algorithm. All the sonic material for this piece is generated by the AI, however, this is not to say that there is no involvement from the composer.

Like any instrument, before we play the JukeBox Algorithm, it must be well tuned, which is achieved through a highly specific process of training.

Unlike the tuning of a traditional instrument, with JukeBox we are not afforded the luxury of relating the pitches of our instrument to any defined reference. Here, the words training and tuning are used interchangeably to describe a continuous process of generalising the timbre of our instrument towards a musical aesthetic, as defined by our training data.

To describe the fulfilment of these aesthetic goals as reflected in the tuning of our instrument I will be using the term “Well Trained”. However, it should be noted that while the tuning process itself is clearly defined, what can be considered “Well Trained” is purely subjective and up to the discretion of the composer.

To determine if the model is “Well Trained” relies on an interaction between the model and composer, where the training is stopped and the AI is played by sampling. The output of the AI is then judged on its ability to convincingly recreate the “feeling” of an input data set.

For example, an Algorithm that is “Well Trained” trained on Bach’s TWTC, should not sound like Bach exactly when sampled, but it should capture the “feeling” of the pieces, to the point where aspects of the initial training data can be clearly heard in the final output.
As this “feeling” cannot be clearly defined, the final cut off point of the tuning process is subject to the aesthetic agenda of the composer.

In this way the composer must also take on the role of instrumentalist, both ‘tuning’ and ‘playing’ the JukeBox model, in order to produce a certain sonic outcome, as well as overseeing the process as composer, with an overarching musical vision.

THE WELL TRAINED ALGORITHM

Step 1: [Model TWTC]

1a. Initial Training on "The Well Tempered Clavier"

Using Zag's fine-tuning notebook, the standard JukeBox model [Model-J], should be trained solely from a folder containing recordings of Bach's The Well Tempered Clavier. As the model trains the number of training steps will be output to the console. Once the composer feels that the algorithm might be well trained the composer should stop the training process and note the number of training steps.

1b. Sampling [Model-A] for a "well trained" [Model TWTC]

The newly trained model [Model-A] should then be sampled with Zag's sampling notebook, to generate a short amount of audio (Sample A). Upon its completion, the (Sample A) should be evaluated by the composer to determine if the algorithm sounds "in tune".

If, according to the aesthetic agenda of the composer, (Sample A) convincingly captures the "feeling" of the training data. Save the training step count for reference and move on to 1c.

Otherwise, If the (Sample A) is deemed unconvincing, delete the sample repeat steps 1a & 1b until a convincing sample is generated.

1c. Saving [Model-TWTC]

With the generation of a convincing (Sample A), [Model-A] can be considered sufficiently "well trained" and now represents as unique instrument with a new tuning system, distinct from [Model-J]. This Model should be backed-up in a separate folder and saved as [Model-TWTC].

1d. Generating (Sample-TWTC)

Using a new Sampling Notebook, [Model-TWTC] should then be sampled, this time with the "ancestral" sampling mode as to generate 3 separate 5 minute (600 seconds) long audio files. Upon the completion of these files they should be evaluated by the composer and the most "in tune" audio file should be selected. This audio (Sample-TWTC) should be saved to a new folder and will constitute the starting point of The Well Tuned Algorithm.

1e. Priming for Step 2

(Sample-TWTC) should then be downloaded, and processed so as to create a new audio file consisting of the final 30 seconds of (Sample-TWTC). This new sample (Step_2_Primer) will be used to prime the creations of the next set of samples, so as to create a smooth transition between the sections. (Step_2_Primer) should then be uploaded to a new folder called "primer" within the same subfolder containing the original training notebook.

Step 2: [Model TWTC-TWTP]

2a. Re-Training on "The Well Tuned Piano."
In the same training notebook the unsaved [Model-A] should then be fine-tuned further, this time on a different folder containing the recordings of La Monte Young's The Well Tuned Piano. As [Model-A] is being re-tuned, the composer should try to stop the training at a mid-point, usually this occurs after about half the training steps needed to "well train" [Model-TWTC]. Again, note the number of training steps each time the training notebook is stopped.

2b. Sampling [Model-A] for a "double trained" [Model TWTC-TWTP]
This new model [Model-AB] should then be sampled with Zag's sampling notebook, to generate a short amount of audio (Sample AB) . Upon its completion, (Sample AB) should be evaluated by the composer to determine if the algorithm sounds "double trained".
If the composer can identify in (Sample AB) aspects of both and The Well Tuned Piano, save the number of training steps and move on to step 2c.
If, according to the composer, (Sample AB) still feels too much like the initial training data then it should be deleted and steps 2a and 2b should be repeated to further train [Model-AB] on the new data set.
On the other hand, if the composer feels that (Sample AB) now feels too much like the new training data and contains no elements of the original training data, (Sample AB) should be deleted and [Model-AB] should be replaced with a copy of the backed up [Model-TWTC]. This copy of [Model-TWTC] should then be renamed [Model-A] and step 2a should be repeated, this time stopping the training at a lower number of training steps before moving on the step 2b.

2c. Saving [Model-TWTC-TWTP]
With the generation of a "double tuned" (Sample AB), [Model-AB] can be considered sufficiently "well trained" and now represents a unique instrument with a new tuning system, distinct from [Model-A]. This Model should be backed-up in a separate folder and saved as [Model-TWTC-TWTP].

2d. Generating (Sample-TWTC-TWTP)
Using a new Sampling Notebook, [Model-TWTC-TWTP] should then be sampled, this time with the "primed" sampling mode as to generate 3 audio files, that are primed for 30 seconds on (Step_2_Primer) and are 5 and a half minutes (630 seconds) long in total. Upon the completion of these samples they should be evaluated by the composer and the most "in tune" audio file should be selected. This audio (Sample-TWTC-TWTP) should be saved to a folder with (Sample-TWTC) and will constitute the middle Section of The Well Tuned Algorithm.

2e. Priming for Step 3
(Sample-TWTC-TWTP) should then be downloaded, and processed in the same way as

in step 1e. This new sample (Step_3_Primer) can then be added to the primer folder, ready to prime the audio files in Step 3.

Step 3 [Model-TWTP]

3a. Well-Training on "The Well Tuned Piano."
In the same training notebook the now "double trained" [Model-AB] should continue the fine-tuning process on the La Monte Young dataset. This time the composer should try to stop the training at around double the steps of the initial [Model-TWTC] training. Here you no longer need to note the training steps.

3b. Sampling [Model-B] for a "well-trained" [Model TWTP]
After this third round of training the new [Model-B] should then used to generate a sample once more. Upon completion, the sample should evaluated by the composer to determine if the algorithm sounds "well trained".
If, when evaluated, the generated sample has lost all of the Bachian characteristics of [Model-A] and now produces a purely "well trained" Youngian (Sample-B) move on to step 3c.
Otherwise, delete sample (Sample-B) and repeat steps 3a and 3b until a "well trained" (Sample-B) is produced.

3c. Saving [Model-TWTP]
With the generation of a "well tuned" (Sample B), [Model-B] can be considered sufficiently "well trained" and now represents a unique instrument with a new tuning system, distinct from [Model-A] and [Model-AB]. This new [Model-B] should be backed-up in a separate folder and saved as [Model-TWTP].

3d. Generating (Sample-TWTP)
Using a new Sampling Notebook, [Model-TWTP] should then be sampled, again using the "primed" sampling mode to generate 3 more audio files 5 and a half minutes (630 seconds) long. this time making sure that the files are primed on (Step_3_Primer). Upon the completion of these samples they should be evaluated by the composer and the most "in tune" audio file should be selected. This audio (Sample-TWTP) should be saved to a folder with (Sample-TWTC) and (Sample-TWTC-TWTP). (Sample-TWTP) will constitute the final section of The Well Tuned Algorithm.

Step 4: Arranging
Finally the generated audio files (Sample-TWTC), (Sample-TWTC-TWTP) and (Sample-TWTP) should be download and imported into a DAW. Here, the first 30 seconds of (Sample-TWTC-TWTP) and (Sample-TWTP) can be removed and the files should be arranged chronologically in order of creation. If done correctly the primed sampling of (Sample-TWTC-TWTP) and (Sample-TWTP) should serve to create a transition between each of the audio files, allowing for the creation of a 15 minute piece of music.

This can then be exported as a single audio file (The Well Tuned Algorithm)
If done correctly, (The Well Tuned Algorithm) should alternate between the different tuning systems as the piece progresses, moving from well defined, equally intervaled notes, characteristic of Gould's rendition of the 48's in the first section into to a combination of frequencies and timbres from both TWTC and TWTP in the second section before finally evolving into the rich justly intoned chordal clouds implemented by Young in his playing of TWTP as the piece closes out.

For a graphical representation of this score see appendix.

CONCLUSION

This project sought to utilise 'instrumentisation' as a means to reconceptualise AI technology as a musical instrument in its own right, with the goal of promoting artistic expression within a greater creative process. Through the lens of this framework, I aimed to identify and utilise the specific musical affordances of JukeBox to compose a musical work, *The Well Trained Algorithm*. Here, I will evaluate the success of this endeavour in relation to the proposed research question.

In the composition of *The Well Trained Algorithm*, I found that through 'instrumentisation' I was able to successfully reconceptualise JukeBox as a musical instrument in order to better utilise the AI as a tool for musical expression. Just as Bach and Young used physical instruments to explore the possibilities of their medium, when utilised effectively through the same lens, an AI system is capable of creating unique and interesting musical sounds. By utilising 'instrumentisation' to integrate AI technology into the creative process, I was able to explore new avenues of expression and push the boundaries of what is possible in music creation. Ultimately, I feel that the use of AI technology as a musical instrument here resulted in an innovative and compelling musical experience that has the potential to be both technologically impressive and artistically meaningful.

I feel that 'instrumentisation' as a conceptual framework was critical in applying a cohesive musical language to this piece. Through my application of this method of reconceptualisation, I was able to identify the specific instrumental affordances of JukeBox, and in doing so, I was able to utilise the technology more effectively to communicate musical ideas within the piece. This method of 'instrumentisation' also helped me to approach the AI within a larger musical context and, in doing so, identify its unique strengths as compared to traditional instruments. This allowed me to rethink my use of this technology, focusing on its input rather than its output, as a means to facilitate more meaningful artistic expression.

I observed that by applying musical terms, as allowed by this framework, I was able to establish a compositional foundation for the piece. Drawing an analogy between the concept of training the AI and a tuning system allowed me to create a parallel between this work and the works of Bach and Young. By reimagining this affordance of the AI as a method of composition in its own right I was able to offer an alternative system for tuning, that bypassed the limitations imposed by traditional instruments. Additionally, through this parallel I was able to situate the work within a larger musical context. This helped to inform the compositional process, delineating separate movements and creating a sense of coherence throughout the piece. By doing so, I established a larger musical structure that demonstrated the unique affordances of the AI instrument, such as "primed sampling" and "double tuning" while emphasising *The Well Trained Algorithm* as a continuation of concepts explored by Bach and Young.

Despite the usefulness of ‘instrumentisation’ to the realisation of this piece, there were some limitations to this model of thinking. In its analysis of JukeBox as a musical instrument, TWTA focuses on one aspect of the use of AI in music, and as such, there were many other aspects that were not explored. This includes the many ethical implications of using AI as an instrument. For example, this piece relies on pre-recorded works of others as a basis for the training data, which could potentially bring into question not only the originality of this work but also the authorship, as currently US law states intellectual property is eligible for copyright protection only if it is the result of human creative effort. Additionally, the training data was limited to two musical works, so any generalisation of these findings to other data sets may be limited. Since many of the processes utilised in this piece rely on the context of the training data, how would the piece be affected if, for example, a different set of Bach recordings were used with a more traditional play style?

The decision to utilise AI for this piece also resulted in several challenges during the composition process. The main problem I faced while interfacing with the JukeBox model is that the process is not very intuitive. There was a steep learning curve, and notebooks required careful execution and system management in order to run correctly and achieve the desired results. As the process relies on a lot of dependencies working in tandem, a small change to any one of these systems would break the entire process, and debugging sometimes took days.

Additionally, since the AI is a black box system, there were challenges in balancing the amount of training in each section so that each data set could be heard clearly without overtraining to the point where the output was too close to the input data set. This resulted in a lot of trial and error and time spent retraining from a previous backup of the model.

In closing, I feel that the reconceptualisation of AI technology as a musical instrument in the composition of The Well Trained Algorithm has proven to be a successful experiment. Through the "instrumentisation", the specific musical affordances of JukeBox were identified and utilised in creating a unique musical experience. The application of this framework allowed for the foundation of a compositional foundation that drew on the parallels between training the AI and tuning a physical instrument and successfully established itself as a development of the work of Bach and Young, and the resulting musical piece demonstrated the unique strengths of AI as an instrument. However, it is important to acknowledge the limitations of this model of thinking, the time-consuming nature of this work, and the technological challenges faced during the composition process. Overall, I feel this piece has opened up new avenues for exploration in music creation and has the potential to promote further artistic expression and innovation in this field that challenges the current perceptions of the use of AI music.

REFERENCES

Allami, K, and Counterpoint, 'Tuning system: La Monte Young's The Well-Tuned Piano (1964)', Leimma, 2021 https://isartum.net/leimma/28/refpitch/Ds3-51/tuningsystem/1r1_1r1~279r256~9r8~147r128~21r16~93r64~189r128~3r2~49r32~7r4~31r16~63r32/, (accessed 7 June 2023).

Arkoudas, K., and Bringsjord, S., 'Philosophical Foundations' in Frankish and W. Ramsey (Eds.), *The Cambridge Handbook of Artificial Intelligence*, Cambridge University Press, 2014, pp. 34-63.

Stephen, B., 'OpenAI Introduces Jukebox, a New AI Model That Generates Genre- Specific Music', *The Verge*, Vox Media, 2020, https://www.theverge.com/2020/4/30/21243038/openai-jukebox-model-raw-audio-lyrics-ai-generated-copyright, (accessed 7 June 2023).

Aswad, J., 'Universal Music Asks Streaming Services to Block AI Access to Its Songs', *Variety*, Variety, 2023, https://variety.com/2023/music/news/universal-music-streaming-services-block-ai-1235582612/, (accessed 25 May 2023).

Bach, J. S., and Palmer, W. A., *The Well-Tempered Clavier, Volume I*, Van Nuys. CA, Alfred Music, 2005.

Bach, J. S. and Novack, S., *The Well-Tempered Clavier: Books I and II, Complete*, Dover Publications, New York. NY, 1983.

Bennett, D., 'AI Deep Fake of the Pope's Puffy Coat Shows the Power of the Human Mind', *Bloomberg*, Bloomberg, 2023, https://www.bloomberg.com/news/newsletters/2023-04-06/pope-francis-white-puffer-coat-ai-image-sparks-deep-fake-concerns, (accessed 5 June 2023).

Campbell, M., Greated, C. A., and Myers, A., *Musical instruments: history, technology, and performance of instruments of western musi*c, Oxford University Press, 2004

Carter, E., 'The Pope Francis Puffer Photo Was Real in Our Hearts', *GQ*, Condé Nast, 2023, https://www.gq.com/story/pope-puffer-jacket-midjourney-ai-meme, (accessed 07 June 2023).

Copeland, B. J., "artificial intelligence". *Encyclopedia Britannica*, 1998, https://www.britannica.com/technology/artificial-intelligence/, (accessed 5 June 2023).

Cope, D., David Cope: 'You pushed the button and out came hundreds and thousands of sonatas' [interviewed by T. Adams], 11 July 2010, *The Guardian*, The Observer, 2010, https://www.theguardian.com/technology/2010/jul/11/david-cope-computer-composer, (accessed 6 June 2023).

Dhariwal, P., et al., 'Jukebox', *OpenAI*, San Francisco. CA, 2020, https://openai.com/research/jukebox, (accessed 5 June, 2023).

Dhariwal, P., et al., 'Jukebox: A Generative Model for Music', *ArXiv,* vol. abs/2005.00341, 2020, http://arxiv.org/abs/2005.00341, (accessed 20 April 2023).

Fridman, L., '#53 Noam Chomsky: Language, Cognition, and Deep Learning', *Lex Fridman Podcast,* [podcast/online video], Presenter L. Fridman, interview with Noam Chomsky, U.S, Lex Fridman, 2019, https://www.youtube.com/watch?v=cMscNuSUy0I, (accessed 5 June 2023).

Gann, K., "La Monte Young's The Well-Tuned Piano." *Kyle Gann*, 1997 https://www.kylegann.com/wtp.html, (accessed 7 June 2023).

Gann, K., "La Monte Young's The Well-Tuned Piano", *Perspectives of New Music*, vol. 31, no. 1 (Winter, 1993), pp. 134-162, https://www.jstor.org/stable/833045/, (accessed 7 June 2023).

Gibson, J., *The Ecological Approach to Visual Perception*, New York. NY, Psychology Press, Taylor & Francis Group, LLC, 1986.

Hanslick, E., *On the musically beautiful,* Translated by Payzant, G., Cambridge. MA, Hackett Publishing, 1986.

Héroux, I., 'Creative processes in the shaping of a musical interpretation: a study of nine professional musicians' *Front Psychol*, vol. 9, no. 665, 2018, https://www.frontiersin.org/articles/10.3389/fpsyg.2018.00665, (accessed 5 June 2023).

Héroux, I., and Fortier, M., *Expérimentation d'une Nouvelle Méthodologie Pour Expliciter Le Processus de Création d'une Interprétation Musicale*, Les Cahiers de La Société Québécoise de Recherche En Musique, 2015, pp. 67–79.

Hofstadter, D. R., *Gödel, Escher, Bach*, New York. NY, Basic Books Inc, 1979.

Hogan, M., "La Monte Young: The Well-Tuned Piano 81 X 25 6:17:50-11:18:59 PM NYC.", *Pitchfork*, Condé Nast, 2010, https://pitchfork.com/reviews/albums/la-monte-young-the-well-tuned-piano-81-x-25-61750-111859-pm-nyc/, (accessed April 21 2023).

Iyanda, M., 'ChatGPT: How the World Is Reenacting Moral Panics on Artificial Intelligence', *Tekedia,* Tekedia, 2023, https://www.tekedia.com/chatgpt-how-the-world-is-reenacting-moral-panics-on-artificial-intelligence/, (accessed 7 June 2023).

Johnston, B., "Pinned Down by the Piano", *The Village Voice,* Village Voice, 2002, https://www.villagevoice.com/2002/09/03/pinned-down-by-the-piano/, (accessed 7 June 2023).

Katz, Y., 'Noam Chomsky on Where Artificial Intelligence Went Wrong', *The Atlantic,* The Atlantic, 2012, https://www.theatlantic.com/technology/archive/2012/11/noam-chomsky-on-where-artificial-intelligence-went-wrong/261637/, (accessed 5 June 2023).

Lerman, R., "At the Colorado State Fair, visitors get a glimpse of A.I.'s mid-journey.", *The Washington Post*, The Washington Post, 2022, https://www.washingtonpost.com/technology/2022/09/02/midjourney-artificial-intelligence-state-fair-colorado/, (accessed 5 June 2023).

Luka, A., *Phainesthai, Reverse-Engineering Auditory Processes as a Tool for Musical Composition,* MA diss., Stockholm, Royal College of Music in Stockholm, 2023.

Malone, K., *Harmonic Space & Hegemonic Process*, MA diss., Stockholm, Royal College of Music in Stockholm, 2017, http://urn.kb.se/resolve?urn=urn:nbn:se:kmh:diva-2416, (accessed 7 June 2023).

Manyika, J., and Bughin, J., ‘The coming of AI Spring’, *McKinsey Global Institute*, Project Syndicate, 2019, https://www.mckinsey.com/mgi/overview/in-the-news/the-coming-of-ai-spring, (accessed 5 June 2023).

Manovich, L, ‘Defining AI Arts: Three Proposals’, *AI and Dialog of Cultures*, [exhibition catalog], Hermitage Museum, Saint-Petersburg, 2019, http://manovich.net/index.php/projects/defining-ai-arts-three-proposals, (accessed 7 June 2023).

McCarthy, J., AI@50 Conference, *What Was Expected, What We Did, and AI Today*, Filene Auditorium, Moore Hall, Dartmouth College, Hanover, NH, July 13, 2006, cited in conference notes by Maker, M. H., (2006) AI@50: AI Past, Present, Future. Archived at https://web.archive.org/web/20081008120238/http://www.engagingexperience.com/2006/07/ai50_ai_past_pr.html on archive.org, (archived from the original on January 3, 2006), (retrieved October 16, 2008), (accessed 5 June 2023).

Meyer, L. B., *Emotion and Meaning in Music*, Chicago University Press, 1956.

Riley, T., Wolf, D., Young, L.M., Zazeela, M., and Farneth, D., *The Well-Tuned Piano 81 X 25 6:17:50 - 11:18:59 PM NYC* [LP liner notes], Gramavision Records, New York. NY, 1987.

Roose, K, ‘An A.I.-Generated Picture Won an Art Prize. Artists Aren’t Happy.’, *The New York Times*, The New York Times Company, 2022, https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html, (accessed 20 April 2023).

Russell, S. J., and Norvig, P., *Artificial Intelligence: A Modern Approach*, Pearson Series in Artificial Intelligence, 4th US edn., Hoboken, Pearson Education Limited, 2022.

Searle, J., "Minds, brains, and programs." *Behavioral and Brain Sciences*, vol. 3 no. 3, 1980, pp. 417-424.

Turing, A. M., ‘I.—COMPUTING MACHINERY AND INTELLIGENCE’, *Mind*, vol. 59, no. 236, 1950, pp. 433–460, https://academic.oup.com/mind/article/LIX/236/433/986238, (accessed 5 June 2023).

Turing, A. M. , 'Intelligent Machinery (1948)', in Copeland, B. J. (ed.), *The Essential Turing*, Oxford, 2004, online edn., Oxford Academic, 2020, pp. 395–432., https://academic.oup.com/book/42030/chapter-abstract/355746030, (accessed 7 June 2023).

Xen Wiki, *Young-lm piano* [website], https://en.xen.wiki/w/Young-lm_piano, (accessed 7 June 2023).

Zhang, D., et al., *The AI Index 2021 Annual Report*, AI Index Steering Committee, Human-Centered AI Institute, CA, Stanford University, Stanford, 2021, p. 2, https://aiindex.org/2021-report/, (accessed 5 June 2023).

CODE

Dhariwal, P., et al., ‘Jukebox: A Generative Model for Music’ 2020, https://github.com/openai/jukebox, (accessed 7 June 2023).

Zag, ‘Zags_1bl_finetuning_l2_labeled_with_lyrics___ft5.ipynb’ [source code], 2022, https://colab.research.google.com/drive/17Uk7R2nG2I5W-unOX2uU576wGdfvI_13 (accessed 6 June 2023). *

Zag, ‘Zags_5_0_ft4_multi17_release.ipynb [source code], 2021, https://colab.research.google.com/drive/1p8VdcnFWofabSom_ILpmpxbhf1JHI-Cq, (accessed 6 June 2023). **

MUSIC

Bach, Johan Sebastian (Comp.), Gould, Glenn,
The Well-Tempered Clavier, Book 1/ Preludes And Fugues 1-8, The Well-Tempered Clavier, 6 vols, vol. 1, Columbia Masterworks, Columbia Records, US, 1963.
, *The Well-Tempered Clavier, Book 1/ Preludes And Fugues 9-16,* The Well-Tempered Clavier, 6 vols vol. 2, Columbia Masterworks, Columbia Records, US, 1964.
, *The Well-Tempered Clavier, Book 1/ Preludes And Fugues 17-24*, The Well-Tempered Clavier, 6 vols, vol. 3, Columbia Masterworks, Columbia Records, US, 1965.
, *The Well-Tempered Clavier, Book 2/ Preludes And Fugues 1-8*, The Well-Tempered Clavier, 6 vols, vol. 4, Columbia Masterworks, Columbia Records, US, 1968.
, *The Well-Tempered Clavier, Book 2/ Preludes And Fugues 9-16*, The Well-Tempered Clavier, 6 vols, vol. 5, Columbia Masterworks, Columbia Records, US, 1970.
, *The Well-Tempered Clavier, Book 2/ Preludes And Fugues 17-24*, The Well-Tempered Clavier, 6 vols, vol. 6, Columbia Masterworks, Columbia Records, US, 1971.

Young, La Monte,
The Well-Tuned Piano 81 X 25 6:17:50 - 11:18:59 PM NYC, Grammavision, US, 1987.

* **Community source code is available to download at Jukebox Community Discord channel which can be accessed by using invite link:

https://discord.com/invite/aEqXFN9amV

Once you have joined the Jukebox Community Discord channel, code can be found in the ‘latest-releases’ channel here:

https://discord.com/channels/766622617393430559/821756207643099216

APPENDIX

The Well Trained Algorithm Graphical Score

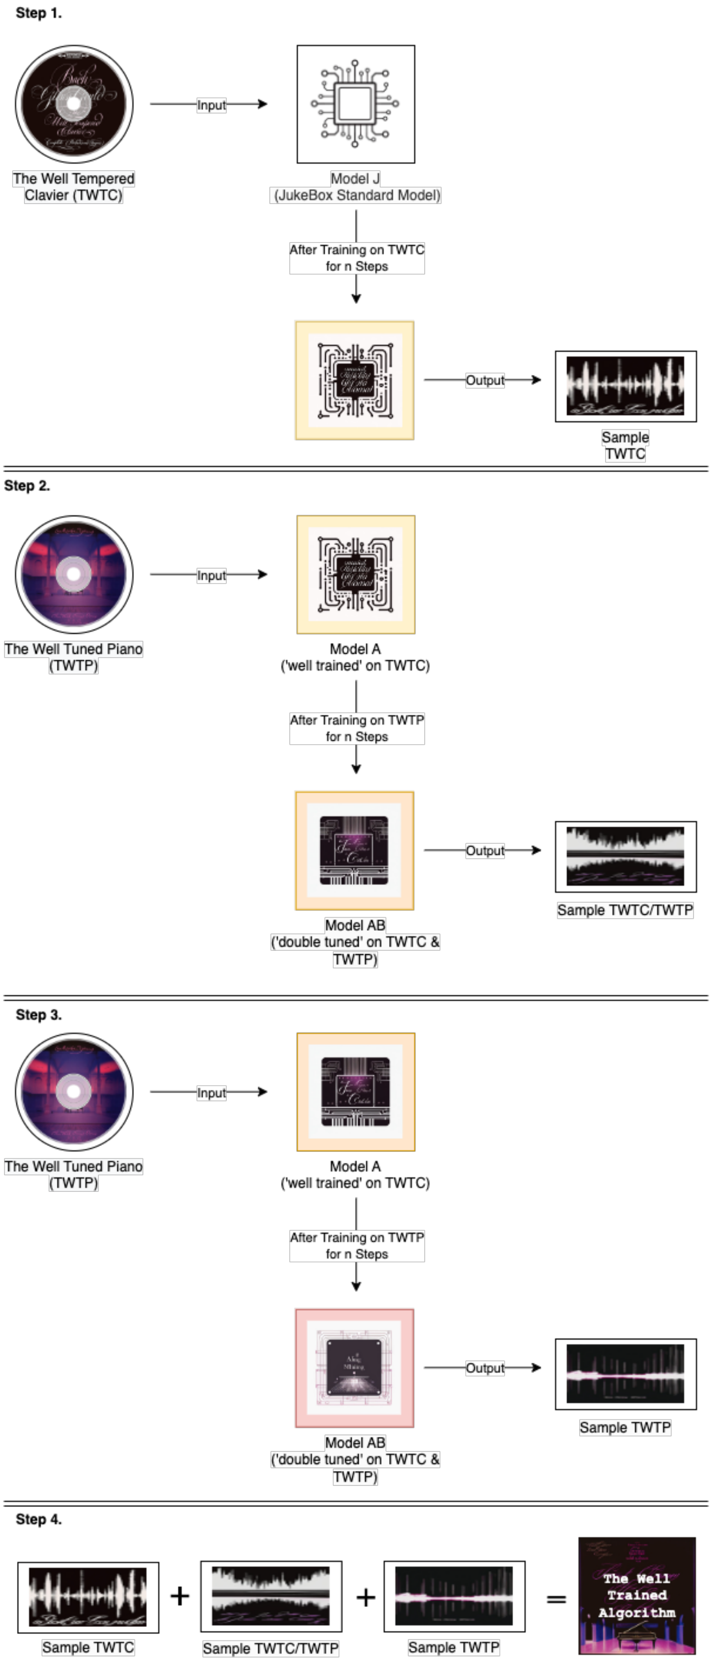

Kungl.
Musikhögskolan