xCoAx 2020
Computation,
Communication,
Aesthetics & X
2020.xCoAx.org
Graz, Austria

Unboxing the Machine:
Artificial Agents in Music

Antonio Pošćić
antonio.poscic@gmail.com
Unaffiliated, Zagreb, Croatia

Gordan Kreković
gordan.krekovic@visagetechnologies.com
Visage Technologies, Zagreb, Croatia

Keywords: Artificial Intelligence, Generative Art, Live Coding, Generative Music, Computer Art.

The constant evolution of philosophical views on arts is interwoven with the trajectories of accelerating technological amelioration. In the current prominent emergence of generative algorithms there is an immediate need for making sense of modern technologies that more and more seem to step in the realm that has been reserved for humans—creativity. This paper aims at demystifying implications of black-box generative algorithms by: 1) depicting the current state of practice and research in this interdisciplinary field, 2) dissecting and examining the constitutional characteristics of artificial agents: artificialness and intelligence, and 3) applying the resulting implicit theory to a practical case of generating a live coding musical performance.

1. Introduction

In the broadest sense, artificial intelligence is any sort of intelligence exhibited by machines (Nilsson 1998). More specifically, the field of computer science that deals with artificial intelligence is commonly concerned with computational agents capable of perceiving their surrounding contexts and undertaking action in an effort to maximize their set goals and expected results (Poole et al. 1998, Russel and Norvig 2003). While this sort of categorization is technologically and philosophically correct, the definition is vague and provides only a basic framework for understanding either the algorithmic mechanisms it comprises or the sociological implications it poses.

With the rise of computational power and ubiquitous computing made available through smartphones and other personal devices, various forms of artificial intelligence have made their ways into all layers of society and daily life (Mlynář et al. 2018). And while the technological progress of the associated techniques and processes is fast and steep, individuals, communities, and societies struggle with understanding their nature and potential ethical pitfalls. Instead, there is often polarization: techno optimists on one side and Luddites on the other, devolving a complex subject into improper analogues and wrong assumptions (Elliott 2014).

As with other aspects of contemporary life, artificial intelligence and machine learning infiltrated artistic practices through osmosis, both in academic contexts and mainstream art (Parker 2019). The way in which artists – and computer scientists turned artists – approach this set of technologies is multifaceted, from simple (and contested) black box usages to complex modifications and subversions of AI mechanisms. Similar to the impact in wider sociological narratives, the usage of AI in art is a controversial topic, more so than any other digital tool has been in the past. Again an axis of opponents and proponents forms, with the first group believing AI to bring on "the death of art" through replacement of true creativity and human artistry, while the latter see immense possibility in the augmentation of human cognitive and performative capabilities.

Taking into account all these premises, we defined the two primary concerns of this paper. The first is to provide an initial demystification and delineation of the various aesthetic and artistic implications of AI in art (visual arts and music, specifically), while also describing the current state of the interdisciplinary field. Here we also present a novel dissection of AI in art, defining its artificialness and intelligence as crucial aspects. The second concern of the paper is to demonstrate an artwork—a participative audiovisual performance—based on generative live coding that is simultaneously born from the basis of well-known AI techniques, but which also inherently problematizes the various dynamics between humans and algorithms. The discussion that led to specific choices in the preparation of the performance thus becomes a part of the artwork itself and serves as a

contribution to the overarching discussion about the position and meaning of artificial intelligence in art.

1.1. Philosophical, Ethical, and Socioeconomic Questions

While this paper focuses on the purely artistic aspects of artificial intelligence, specifically in music, we recognize that this domain is not isolated and exists within a larger social framework. In this sense, while beyond the scope of this paper, we identify that there are philosophical, ethical, and socioeconomic issues with the use of AI and machine learning for artistic purposes.

When philosophical questions are concerned, the notion of intelligence itself is loaded, with the approaches to defining and categorizing it a matter of many researches and polemics (Sternberg 2003). In the context of art, the main philosophical question lands somewhere between ontology, semiotics, and aesthetics. Can a computer replace a human in creative tasks? It is an open-ended question with unclear conclusions. Currently, the prevailing view is reductive in that computational creativity cannot create original art in the same way that humans do (Magnusson 2019). Still, the discussion is ongoing and productive, walking in step with developments in technology.

The second type of issues with AI in art is related to ethical dimensions of its applications. For example, questions of copyright and intellectual property come into light, with AI systems like generative adversarial network (GAN) introducing the capability of autonomously generating new art based on patterns learned from existing data sets (McCormack 2014). Similarly, the employment of AI in music distribution can shape and drive musical trends and influence long-term development of the arts (Kaplan and Haenlein 2019).

This final point is also related to socioeconomic questions hovering around AI. Mainly, with GANs generating a limitless number of compositions in the style of Muzak (White and Matulionyte 2019), they reduce the space for human composers to work in profitable fields. While this scarcely affects the truly creative branches of art (see second paragraph in this chapter), many musicians use these jobs as source of income and to finance their main artistic outlets.

2. Artificial Intelligence in Art

If ideas of using generative and proto-artificial intelligence techniques in art have been around for several decades (Kugel 1981, Galanter 2003), their wider and consistently applied use is still a fresh and relatively unexplored field. Largely owing to the increase in computational power available to individuals both in mainstream and academic settings, advances in AI algorithms, and open source availability of these tools, the latter part of the 2010s saw a significant increase in attempts to use AI for artistic purposes. The domains in which this phenomenon manifested are many, but

the progress is fastest and most easily observable in visual arts and music, which follow similar tracks and on which similar concepts can be applied through mathematically described signals. In the next chapters we give an overview of some significant recent works and modes of use of AI in the aforementioned fields.

2.1. Visual Arts

Championed by artists like Mario Klingemann, Memo Akten, Robbie Barrat, Gene Kogan, and Mike Tyke, generative adversarial networks have been the most widely accepted form of AI in visual arts (Schmitt 2018). These artists use publicly available algorithms and tools, train them on various data sets of existing art—masters' portraits, for example—modify their parameters, and have them generate new pieces in a style that mimics the original data sets, but which also introduces or allows surfacing specific artifacts inherent to the used technologies. The resulting aesthetic is one that is still very much human, not machinic, yet which reveals some of the algorithm's intrinsic properties. Here, the artist has three roles: to select the data sets used for training the system, to adjust the parameters of the system, and to finally act as a curator who selects the most compelling pieces in a vast space of generated works.

Apart from GANs, artists employ various evolutionary algorithms to generate dynamic 3D artworks (Romero 2008), virtual reality pieces (Lugrin et al. 2006), and abstract works whose aesthetic could be described as mathematic and detached from human preference (Wannarumon et al. 2008). Throughout most of these approaches, and as we will explain later, AI is used within the boundaries of the technology's original parameters and most often as mere tools or black boxes. Even if we could state that for fields or artistic approaches still in their infancy every work can be considered questioning of its toolset and lineage, there are artworks which make the problematics of the technology itself their focal point (Roh 2018). While rare, these attempts function on multiple levels: as proponents of the methods employed and as interrogators of the underlying motives and concepts.

2.2. Music

Unlike visual arts where we can identify a common mode of AI use through GANs and adjacent scenes forming, in the context of music AI has been used somewhat sparsely and disparately. Instead of trying to summarize current meta-narratives, we focus on four distinct pieces which illustrate how AI can be used for music creation purposes.

Using a modified SampleRNN architecture, Zukowski and Carr created the Dadabots system and employed it to generate black metal and math rock similes (Zukowski and Carr 2018). While this approach has many similarities with the use of GANs in visual arts, the distinction is that the material

that they use is both contemporary and highly aesthetically recognizable. Additionally, the aesthetic and stylistic elements of black metal and math rock can be considered chaotic and difficult to analyze using conventional techniques due to their use of timbre and space as compositional drivers (Lee et al. 2009). Because of this and the contemporary nature of the originating styles, the pieces produced by Dadabots discover a completely new aesthetic, birthing an acousmatic experience out of isolated and subverted black metal elements. As the authors state, "we are delighted by the unique characteristic artifacts of neural synthesis", emphasizing the different role and expectations of AI in art as contrasted to AI for general purposes. Here, errors are cherished as the researches/musicians curate and select pieces from Dadabots' vast output.

While Zukowski and Carr welcome the unexpected outcomes of their process, researchers and musicians Holly Herndon and Mathew Dryhurst created and trained an AI named Spawn for Herndon's *PROTO* album with a specific role in mind. The process of training and using the AI is iterative in this case, as Herndon created and recorded music, fed it to Spawn, and then finally used the system to transform and generate new samples. These were then arranged into the final compositions. Here the AI, while fairly complex and well designed, is used as a mere tool for generating sound samples, just another "voice singing in unison", while the final act of composing is in Herndon's hands. This resulted in a record that, despite its advanced evolution, fits within expectations of conventional electronic music. Like with Dadabots, there is an additional dimension of performativity and curious exploration of the system itself at play, as Herndon mentions naming, anthropomorphizing, and raising the AI entity.

If Herndon and Zukowski and Carr reside on opposite sides of a continuum, experimental and drone music duo Emptyset fall between them, opting neither for a fully generative process nor choosing to harness the AI purely as a sample-generator. James Ginzburg and Paul Purgas instead employed machine learning techniques to explore new, unexpected possibilities out of data sets they sourced themselves. They explain their artistic process behind the record Blossoms and departure towards a new aesthetic framework:

> "The machine learning system for Blossoms was developed through extensive audio training, a process of seeding a software model with a sonic knowledge base of material to learn and predict from. This was supplied from a collection of their existing material as well as 10 hours of improvised recordings using wood, metal and drum skins. This collection of electronic and acoustic sounds formed unexpected outcomes as the system sought out coherence from within this vastly diverse source material, attempting to form a logic from within the contradictions of the sonic data set. The system demonstrates obscure mechanisms of relational reasoning and pattern recognition, finding

correlations and connections between seemingly unrelated sounds and manifesting an emergent non-human musicality.” (Smart 2019)

Finally, a piece that focuses on the phenomenology of AI itself is composer Jennifer Walshe’s work *ULTRACHUNK*, realized in collaboration with visual artist and researcher Memo Akten, which explores the emerging world of computational intelligence. *ULTRACHUNK* is an improvised piece, a duet between Walshe and an AI which acts as a mimetic partner and absorbs the main characteristics of Walshe’s identity, namely her voice and face.

As will be investigated in detail in later chapters, none of these systems demonstrate real intelligence or creativity, which ultimately still reside within the human authors and operators of the used systems. They are not magical, mystical, or dangerous contraptions, but rather tools in the hands and minds of creative organic artists.

2.3. Mystified Black Boxes, Purposeful Tools, and Cognizant Systems

“Instruments with machine learning capabilities will learn from their players, who ‘train’ them to adapt to their playing, so no instrument will be the same.

The instrument often becomes a piece in itself, as the creator of the instrument has some musical purpose in mind, often quite specific. The boundary between a piece and an instrument is deliberately vague, and it can fluctuate.” (Magnusson 2019)

In the above quote, the author speculates on a future where the AI in music is truly intelligent and creative, capable of higher cognitive functions compared to the current baseline. This scenario envisages a future in which AI instruments are omnipresent and infinitely varied, tailored by and to their users. In contrast to these hypothetical cognizant systems, the methods employed by musicians active today, and illustrated through the examples in the above chapters, are still rudimentary.

An often encountered approach is to consider the AI system to be black boxes, impervious to the musician’s full understanding. In this case, the musician inputs some data into the system and collects the result, integrating it into their pre-existing aesthetics. An analogy in the world of physical instruments would be a pianist that only uses the keyboard of a piano and resorts to traditional, pre-20th century modes of playing.

Opposed to this basic approach are artists who understand and modify the behavior of the system. Most of the examples detailed above fall into this category and make the artist not only the final user of the system, but an instrument builder. Mario Klingemann and Holly Herndon, for example, both understand the concept and functionality of the AI they’ve either built or use, tweaking parameters and influencing its behavior on deep levels.

In the analogy from the world of physical instruments, they are akin to pianists who open the bodies of their pianos, tug at strings, prepare the instrument with objects, etc. It suggests a deeper and, in a way, subversive understanding of the technology in play.

Returning to Magnusson's concept, there are no current examples of an AI employed by artists that would be completely and autonomously tailored to them. To understand what this jump—from black boxes and ready-made instruments to cognizant systems—would entail, in the next chapter we dissect artificial intelligence in art through dimensions of artificialness and intelligence.

3. Artificialness and Intelligence

Increasingly perceived as a modality of AI that has a formidable potential of disrupting arts by challenging the most intrinsic and delicate questions about uniqueness and human creativity, generative algorithms have become a trending controversial topic on the cross section of arts, technology and philosophy (Nake 1971; Schneider & Rea 2018). It may feel like artificially intelligent creativity is establishing a threatening position from which it can easily endanger the pure meaning and essential values of arts that have been developed through centuries of continuous cultural tradition and incremental advancements. However, we argue that such generative AI can be treated as modern brushes and palettes that serve as tools for artistic creation opening unexpected possibilities and novel modalities inherent to the actual technological and social state of progress. Those tools indeed change the value system of art, but this system has been always in a continuous evolution following changes in the civilizational context. The pace of evolution now seems significantly accelerated and the urgency of finding extended narratives about the position, meaning, and appreciative dimensions of art in the age of artificial intelligence is alarmingly rising.

The connection between brushes, palettes, and various generative algorithms holds in nominal characteristics of the latter: artificialness and intelligence. Observing these characteristics through development of technology reveals the backbone of continuity which may also serve for constructing desperately needed new narratives.

3.1. Artificialness

Artificialness, defined as a condition of lacking naturalness or spontaneity, in the context of traditional and more recent (semi)autonomous tools may be considered in terms of quality and quantity of human involvement in making, adjusting, supporting, and using tools to reach desired results. Being very intuitive (from the current point of view), early tools like brushes and palettes may be observed as natural and thereby non-artificial. Indeed, paint is one of the earliest inventions of mankind that is likely to be approximately

100,000 years old (Mayell 2004). However, even such a rudimentary tool made a disruptive impact, at least in the sense of long-term preservation of human artifacts that allowed modern archeologists to appreciate and study them. Individuals who possessed paint had an opportunity to express themselves in a specific way that was not available otherwise, and an opportunity to leave a trace. As a result of a human intervention of intentionally applying pigments to leave visible colorful traces, paint can be considered as a tool that was, in line with the time of its creation, somewhat artificial and disruptive. A possibility of painting opened numerous approaches and styles and led to a complex value system related to visual arts.

Along with iterative improvements, some technological advancements triggered a much faster shift of artistic focus and expansion of the value system. The emergence of photography in the 19th century trivialized the problem of applying color on canvas to create realistic images. The process for achieving this goal translated from the painting skills to intrinsic mechanisms of the tool. Photography could have been seen as a threat to classical painting in a way that GANs may be considered as the same threat nowadays. However, there are at least two reasons why brushes and palettes survived the 19th century. The first one is that the emergence of the new tool did not prevent painters from exploring novel and creative ways of using old ones. Moreover, photography provoked accelerated changes that led to novel styles. Similarly, when IBM's program Deep Blue won a chess match with Garry Kasparov in 1997, humans did not stop playing chess, but they became even better by learning from the system and trying to defeat it (Harari 2018).

The second reason for the coexistence of brushes, cameras, and GANs is that none of those tools can provide completeness in comparison to the scale at which humans currently understand and perceive the world. While photography resolved the problem of capturing visual reality, artistic ambition was and will always be much more than that. Similarly, GANs may complement photography due to their capability to create novel imaginative and realistic images to a desired extent. Although this type of advancement raises serious challenges to existing forms of visual arts, it is far from being complete. Current GANs excel in creating stunning new artifacts that follow statistical distributions of existing images used for training, but they are still incapable of understanding the social, cultural, economic, political, and even artistic context that influences art produced by humans. For that reason, the role of artists operating GANs and curating results is imminent. Almost without doubt, the future will bring advanced AI that can harvest more data from different sources and thereby incorporate some aspects of the civilizational knowledge and humanic experience in the generated art. However, even then, the artist's role (although significantly altered) would be present as a curator, creator of tools, or at least as a data source. As long as the art is consumed by humans and produced by man-made tools, it will not collapse into itself.

Levels and consequences of artificialness constantly change in the course of time as well as their perception. Modern tools seem to possess the highest possible level of artificialness, but this is due to the tendency of comparing the current state with the known past, while the unknown future may bring much higher levels of artificialness. When humans became used to photos, cameras did not seem so artificial any more. Time and habits *naturalize* artificialness.

Another important observation is that higher levels of artificialness usually mean greater technical complexity, but simpler and more intuitive basic usage. At first, it may seem that modern tools reduce the necessity for fine skills that are traditionally needed to produce art thereby jeopardizing some traditional artistic values. Black-box usage leads to immediate results that eventually become expected and common. On the contrary, operating complex tools to produce more meaningful or authentic results requires mastery, which is proportionally demanding to the complexity of the underlying technology. From that point of view, there is no fear that humans will delegate their creativity to generative algorithms.

3.2. Intelligence

The notion of intelligence has historically relied on implicit theories and definitions constructed from expert opinions (Sternberg 2003). One of the most well-known studies of experts' definitions of intelligence conducted by the editors of the Journal of Educational Psychology (Thorndike 1921) was an early proof of somewhat different, yet overlapping views on human intelligence that further evolved through later attempts of compiling experts' definitions (Sternberg & Detterman 1986). A consensus is equally needed when it comes to computational intelligence. Having more limited expectations on machines than on humans, the understanding of computational intelligence historically shifted following the trajectories of technological improvements (Poole et al. 1998). Computer's ability to efficiently solve numerical and combinatorial problems seemingly significantly overcomes some aspects of the human intelligence, but when it comes to understanding, self-awareness, learning, emotional knowledge, reasoning, planning, creativity, and critical thinking, before recent emergence of deep learning, achievements in computational intelligence were disappointedly modest.

In the context of generative AI in art, we are particularly interested in those qualities that an artificially intelligent system should expose in order to be considered intelligent. When observing an AI system as a black-box tool, it is impossible to assess its computational intelligence form one or just a few outputs. We argue that such a generative intelligence is reflected in the distribution of outputs (that can be statistically approximated using a sufficient number of real outputs) and the understanding of its inner working.

Achieving the satisfactory distribution of outputs is necessary but not sufficient sign of computational intelligence of a generative system. As

classifiers and regressors can be evaluated based on their accuracy, the alignment of distribution of outputs with either implicit or explicit expectations serves a measure of how well generative algorithms solve their generative tasks. Usually, a desirable characteristic of generated material (that also applies to generated art) is to hold a fine balance between known and novel. For example, AI that generates hyper-realistic environments for the simulation purposes is expected to create and equip spaces that are feasible, familiar, and likely to occur in reality, but at the same time authentic, unique, and non-trivially derived from real existing environments. The same is with music generation that aims at finding novel expressions, while relying on or incrementally extending those values of sound organization that underlie the desirable musical consequences.

Inner working of generative algorithms and data necessary to train or feed the algorithm are tightly related to the distribution of its outputs. If two algorithms create outputs with similar statistical distribution, one that relies more on procedural and structural aspects will encapsulate more complexity and require less data than an algorithm with less procedural details, but with better ability to learn. Although not related to generative art, an excellent example is Google's AlphaZero program that in 2017 defeated the almost thousand times more computationally powerful Stockfish 8 program, the world's computer chess champion for 2016. The importance lies in the fact that AlphaZero did not learn to play chess from humans, but from playing it with itself using a technique of reinforcement learning. And it took only four hours. Its radically different inner working led AlphaZero to the victory that proves higher computational intelligence.

While some algorithms can directly compete with each other in a game with explicit rules or they can be compared by their objectively measured performance at solving a given task, such a comparison is more difficult for generative algorithms, especially in the artistic setup. As the valuation of their generated outputs is subjective and they are expected to form a wide distribution (in contrast to optimizing the solution to a problem), computational intelligence is a more delicate notion. Instead of taking into account only the quality of outputs, the inner working of generative algorithms may reveal their ability to produce more or less intelligent outputs. While the discussion at this level remains abstract, more tangible insights into computational intelligence are given in the discussion of the practical part of the study.

4. A Case Study: Generative Live Coding

A complementation to a theoretical view on generative art presented in this paper is a practical exploration of applying AI to live coding. As the art of using computer programming, algorithms, and code as makeshift scores and music creation tools, live coding is built around improvisation, with

musicians most often writing code in real time during live performances. The exclusion of a human performer from such a setup is a radical and metaphoric intervention, as the very purpose of a programming code is to mediate between a human and a machine. Without a human, the machine interfaces only itself exposing an intermediate creative step to potential human observers in a readable form. Autonomous code generation thereby brings a conceptual value, but also entails aesthetic consequences imposed by the live coding environments.

The aim of this practical study was to design and develop a generative system that autonomously creates and sequentially executes blocks of code in TidalCycles, a textual programming language and an extensive library for live coding (McLean and Wiggins 2010). Besides pure code generation, we extended the system to support participative performance and allow audience members to interact with the generated code. The participative influence on the generated music metaphorically represents a modern relation between computers and their users in which users have a limited control, yet a perception of empowerment. The system was intended to generate hours of music following some predefined compositional patterns and aesthetics.

Design decisions in developing the generative algorithm, which is the heart of this artificially intelligent system, emerged from the theoretical discussion and the possible levels of artificialness and intelligence.

By its purpose and functioning, the system is inherently highly artificial, as the aim was to emphasize the artificial creativity and its relations to human actors. Therefore, a possible range of artificialness considered during the system design phase was fairly narrow and mostly reduced to the presentational level where the nuances of artificialness may affect the impression that the system with its generated code and music leaves to observers. One such presentational aspect was the integration of the system within the live coding environment. If we decided to direct the system's outputs to a textual file and execute the generated code blocks manually, this would look less artificial and closer to the way human performers use live coding environments, even though the backend system still remains responsible for the creative part. On the other hand, automatic execution of the generated code emphasizes the impression of the system full autonomy which may be perceived as a higher level of artificialness.

With a growing accumulated assortment of possible generative techniques, the designing phase left a lot of room for adjusting the desired level of the system's intelligence. As the aesthetic and stylistic consistency was a target characteristic, using a sufficient number of predefined blocks of code (either as a training set or atomic building blocks) was a valid way forward. One of the most basic approaches was constructing a Markov chain with predefined code blocks written in its states and transitioning probabilities set manually. In a general case when states are fully connected, the number of transitioning probabilities exhibit quadratic growth making

manual adjustments tedious. An advancement of this method towards a more intelligent system would be implementing an automatic, data-driven approach to setting the probabilities. In a particular implementation, we opted for a formula that makes probabilities of transitioning between two states negatively correlated with the Levenshtein distance between code blocks written in those states.

This solution favors smaller changes from one block to another in the same manner as a human musician during live performances more usually modifies existing code blocks than writes new ones from scratch. However, due to the non-deterministic nature of the Markov chain and further indeterminacies in the formula for calculating transition probabilities, this solution still can generate quite interesting and unexpected sequences of changes. As the third version of a generative algorithm, we could have used a long-short memory (LSTM) neural network that is capable of autonomously generating text by learning from a given training set. This ultimate version would require predefined code blocks just for the training, while its output can be any sequence of characters with the obvious goal to achieve that those characters represent a valid TidalCylces code that preferably produces music in a given stylistic frame.

The three mentioned approaches (basic Markov chain, advanced Markov chain, and the LSTM network) embody the gradation of the system's intelligence determined by its capability to generalize: the approaches based on the Markov chains can generate only preexisting code blocks, while the LSTM is capable of creating blocks not seen in the training set. However, in all these setups, the role of the composer is imminent and prominent. Even with the highest artificial intelligence of the system, the style and the meaning of the output depends on the human creator.

Achieving a fine balance between diversity of the generated material and its aesthetic and structural qualities was the key challenge of developing a system intended to produce hours evolving music which forms a meaningful performance. While the different levels of artificial intelligence set a different contexts for tuning the algorithm's creativity, in all cases of data-driven or machine learning approaches, the resulting artificial creativity depends on the material—building blocks, training data, and parameters of the algorithm. The current AI is far from using a general, civilizational knowledge elicited from global, unstructured and autonomously discovered data sources, so the most important role of human creators is to carefully choose the material. In the case of our generative system, the ability to prepare adequate data also influenced the selection of the algorithm. As LSTM requires large amounts of data that depend on its own complexity, we had to start from simpler algorithms with lower intelligence and lower inherent creativity. However, with almost a thousand of carefully prepared code blocks and the unique semi-autonomous approach of setting transitions between them, we managed to achieve a desired level of diversity while

maintaining a surprise factor as there are theoretically around 10^{32} ways to generate one minute of a performance with, of course, very different occurrence probabilities.

Conclusion

This paper aimed first to demystify, then to dissect, and finally to apply knowledge about the use of artificial intelligence in art, specifically music. To do so, in the first part of the paper we canvassed the current state of the art and examined how various artists employed AI in their work. Here we posited that even behind black-box approaches there is nothing inherently neither dangerous nor problematic when using AI for artistic purposes since, in most cases, the resulting aesthetics and artwork are under control of the human artist. Simultaneously, we identified that in the current state of technology computational agents are incapable of producing truly original and creative art without the involvement of a flesh and blood artist. In other words, AIs become augmentations rather than substitutions.

To better understand these conclusions, we separately approached the nominal components of AI in art: artificialness and intelligence. What does it mean for a generative system to be artificial and intelligent? Through these aspects we described how our current tools behave, posited how they could behave in the future, and explored what this meant for artistic applications. Finally, we discussed this approach and developed considerations of using AI for interactive music performances based on a case study of generative live coding performance.

There are two lines of thought that can be traced throughout the paper. First, there is nothing inherently inexplicable or impenetrably oblique about generative systems and artificial agents. Today, they are mostly used as tools by visual artists and musicians who ultimately impose their own aesthetics and sense of creativity on the machine's output. Secondly, there are unexplored ways of subverting and using AI to further augment creativity in humans by relying on the machine's alien way of rationalizing and "thinking" in order to discover new phenomenologies and aesthetics. This can be further emphasized if the AI is fully controlled or perhaps even built by the artists themselves.

While we avoid speculating about future developments in technology, it is almost certain that artificial intelligence and various generative systems will play an important part in the art of the future. Even if we imagine that these systems might, at some point, eclipse their current role of pure appliances or digital musical instruments and become creative in the narrowest sense of the world, there will still be ample space for true human artistic creativity. Because these systems are, ultimately, mere reflections and extensions of ourselves.

References

Elliott, Stuart W.
2014. "Anticipating a Luddite revival." *Issues in Science and Technology* 30, no. 3: 27-36.

Galanter, Philip.
2003. "What is generative art? Complexity theory as a context for art theory." *GA2003–6th Generative Art Conference*.

Harari, Yuval Noah.
21 Lessons for the 21st Century. *London: Jonathan Cape*. 2018.

Kaplan, Andreas, and Michael Haenlein.
2019. "Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence." *Business Horizons* 62, no. 1: 15-25.

Kugel, Peter.
1981. "Artificial Intelligence and Visual Art." *Leonardo* 14, no. 2: 137-139.

Lee, Chang-Hsing, Jau-Ling Shih, Kun-Ming Yu, and Hwai-San Lin.
2009. "Automatic music genre classification based on modulation spectral analysis of spectral and cepstral features." *IEEE Transactions on Multimedia* 11, no. 4: 670-682.

Lugrin, J., Marc Cavazza, Mark Palmer, and Sean Crooks.
2006. "Artificial intelligence-mediated interaction in virtual reality art." *IEEE intelligent systems* 21, no. 5: 54-62.

Magnusson, Thor.
2019. *Sonic writing: technologies of material, symbolic, and signal inscriptions*. London: Bloomsbury Academic.

Mayell, Hillary.
2004. "Is Bead Find Proof Modern Thought Began in Africa?" *National Geographic News*. Retrieved May 20, 2016.

McCormack, Jon, Oliver Bown, Alan Dorin, Jonathan McCabe, Gordon Monro, and Mitchell Whitelaw.
2014. "Ten questions concerning generative computer art." *Leonardo* 47, no. 2: 135-141.

McLean, Alex, and Geraint Wiggins.
Tidal–pattern language for the live coding of music. In Proceedings of the 7th sound and music computing conference. 2010.

Mlynář, Jakub, Hamed S. Alavi, Himanshu Verma, and Lorenzo Cantoni.
2018. "Towards a sociological conception of artificial intelligence." *International Conference on Artificial General Intelligence*: 130-139.

Nake, Frieder.
There Should Be No Computer-Art. In Bulletin of the Computer Arts Society, London) No. 18, Oct. 1971, 1-2. – Reprinted in: A. Altena, L. van der Velden (eds.): The anthology of computer art. Amsterdam: Sonic Acts 2006, 59-60.

Nilsson, Nils John.
1998. *Artificial intelligence: a new synthesis*. San Francisco, CA: Morgan Kaufmann.

Parker, James R.
2019. *Generative art: algorithms as artistic tool*. Calgary: Durvile Publications.

Poole, David I., Randy G. Goebel, and Alan K. Mackworth.
1998. *Computational intelligence*. New York: Oxford University Press.

Roh, Jinah.
2018. "Artificial Intelligence Art: A Case study on the Artwork An Evolving GAIA." *The Journal of the Korea Contents Association* 18, no. 5: 311-318.

Romero, Juan J.
2008. *The art of artificial evolution: A handbook on evolutionary art and music*. Berlin: Springer Science & Business Media.

Russel, Stuart, and Peter Norvig.
2003. *Artificial Intelligence: A Modern Approach*. London: Pearson Education Limited.

Schmitt, Philipp.
2018. "Augmented imagination: machine learning art as automatism." *Plot (s), the Design Studies Journal* 5: 25-32.

Schneider, Tim & Naomi Rea.
"Has artificial intelligence given us the next great art movement? Experts say slow down, the 'field is in its infancy.'". *Artnet News*. September 2018.

Smart, Dan.
2019. "Experimental duo Emptyset to get all ambitious and flowery on forthcoming AI album (despite the relatively dispassionate band name)." *Tiny Mix Tapes*. https://www.tinymixtapes.com/news/emptyset-announce-share-new-single-petals-forthcoming-ai-album-blossoms

Sternberg, R. J., & Detterman, D. K.
1986. *What is Intelligence? Contemporary Viewpoints on its Nature and Definition*. Norwood, NJ: Ablex.

Sternberg, R. J.
2003. "Intelligence." In *Handbook of psychology: History of Psychology, Vol. 1* edited by D.K. Freedheim, 135-136. New York, NY: John Wiley & Sons Inc.

Thorndike, Edward L.
1921. "Intelligence and its measurement: A symposium." *Journal of Educational Psychology* 12, no. 3: 124.

Wannarumon, Somlak, Erik LJ Bohez, and Kittinan Annanon.
2008. "Aesthetic evolutionary algorithm for fractal-based user-centered jewelry design." *AI EDAM* 22, no. 1: 19-39.

White, Courtney, and Rita Matulionyte.
2019. "Artificial Intelligence Painting The Bigger Picture For Copyright Ownership." Available at SSRN 3498673.

Zukowski, Zack, and C. J. Carr.
2018. "Generating black metal and math rock: Beyond bach, beethoven, and beatles." *arXiv preprint arXiv:1811.06639*.