

Article

Commercial Generative AI as a Tool—The Control–Convenience Spectrum

Krzysztof Cybulski

Faculty of Composition, Theory of Music and Sound Engineering, The Feliks Nowowiejski Academy of Music, 85-075 Bydgoszcz, Poland; k.cybulski@amfn.pl

Abstract

AI-generated content—spanning text, imagery, and music—is becoming increasingly commonplace. As the newest generation of song-producing AI systems garner attention, serious questions emerge regarding the role and place of music producers, particularly in the area of non-artistic, or “utility music”. While it might seem that human skills and creativity are unlikely to be replaced entirely by generative AI in domains such as art music or live performance, recent developments in the field suggest that human efforts in creation of advertisement or background music are already being challenged by generative AI systems. However, there is a number of alternative, more balanced forms of human–machine co-creativity. It is in this regard that I am posing a question: can commercial generative AI systems really be classified as tools in the strict sense of the term? In this paper, I am attempting to answer this question by introducing the “Control–Convenience Spectrum”—a concept I believe applies to all human creative processes that utilize tools. It bears some similarities to earlier ideas in complexity theory or flow psychology—particularly, it proposes that the extremes of this spectrum are unlikely to produce compelling aesthetical outcomes or satisfying creative practice. I argue that prompt-driven commercial generative AI systems occupy one of the far ends of the spectrum, thus failing to meet the criteria for a creative expression tool.

Keywords: generative AI; generative art; generative music; LLM; tool; flow

1. Introduction

While preparing a publication on relations between the creative practice and AI, it might be tempting to focus on describing the latest technological developments in the field, providing historical background for the current technology, or listing examples of the latest AI-produced artwork. Instead, I am attempting a more general overview of the current state of affairs, in order to extract some general rules governing the relations between creative practitioners and technology, be it AI, digital, or any other kind of technology (to the point of considering simple tools, such as paintbrush, the manifestations of technology as well).

Therefore, unconservatively, in the introduction, above all I would like to state what this paper is not.

It is not an attempt at an up-to-date description of the state of the art regarding the AI tools for creative practitioners. The rate of technological advances in the field is so fast that it is difficult to discuss them in detail in the form of a scientific paper, as the production cycle renders almost any attempt at such a description out of date by the time it gets published.



Academic Editor: Ivana Medić

Received: 19 December 2025

Revised: 21 January 2026

Accepted: 31 January 2026

Published: 4 February 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Apart from a few specific examples, the paper is also not a survey of practical implementation of AI in the creative workflows of contemporary artists. Such survey could be easily found elsewhere and has already been extensively covered in a number of publications (Ploin et al. 2022; Moura 2024) and—probably more fittingly considering the fast-pacing development of the medium—in an online form^{1,2}.

I will, however, refer to selected historical facts regarding the development of AI to lay the ground for discussing what currently seems to be a prevailing perception of AI as a novel technology.

Furthermore, it must be stressed that the main portion of this paper regards a specific type of AI-powered systems for artistic creation, which could be characterized as “commercial generative AI”, even though—as I will try to convey in the following sections—they only form a narrow subsection of a broad spectrum of available AI tools. This subsection, however, gained significant publicity in recent years, as opposed to relative obscurity of the majority of other AI tools. For many artists who do not use technology in their daily creative practice, commercial generative AI systems became synonymous with AI in general, with the providers of such systems making efforts to solidify this association.

2. The Control–Convenience Spectrum: Initial Definition

In order to discuss the control–convenience spectrum, we must agree upon basic definitions of “control” and “convenience” terms, as well as what is the general goal of the interaction with generative AI systems.

The concept of the control–convenience spectrum stems from my own experiences as an artist. As a founding member of the panGenerator new media art collective, I have been a part of more than thirty interactive art projects over the last fifteen years, and created another two dozen sound installations and musical instruments on my own, which granted me a sufficient amount of firsthand observations and experiences concerning the relations between the user interface and the user experience. This led me to realize that there always seems to be an inverse proportion between the ease of use and the quality/complexity of the content, generated as a result of the interaction. This inverse proportionality could be visualized as a slider (as in a sliding potentiometer, Figure 1) that controls the proportion between the two variables—ease and control. Therefore, the proposition is to call this inverse proportionality a control–convenience spectrum.



Figure 1. Control–convenience spectrum visualized as a slider.

Interaction design for interactive installations targeted at the general public, as opposed to professional users (e.g., musicians), is obviously a very specific task. Designing a general purpose consumer product is, however, a task rather commonly encountered by designers. On the other hand, designing a musical instrument, meant as a tool for skilled professionals, is easier in certain regards (i.e., some user skills may be taken for granted), but also requires from the designer a deeper knowledge of the discipline for which the tool is going to be used.

The distinction between AI tools for professional users as opposed to entry-level ones is similar. The narrative behind the latter, however, tends to suggest that they are capable of providing the same level of professional output as the former ones. For the sake of this discussion, we will assume that the goal of using the commercial generative AI systems is to obtain artifacts that match the level of professionally produced music or images, regardless of the skill level on the user side. Similarly, when discussing non-AI

tools, we will assume that the goal is not merely producing a single note with a musical instrument or a colorful spot with a paintbrush, but the ability to match the general notion of a “musical performance” or a “painting”.

The “convenience” (or “ease”) term will be used in a similar vein as the “affordance” term, used in the context of design—it will be the measure of how easy it is to interact with the tool at hand, without prior knowledge or experience with the tool and its interface. A highly convenient user interface will be easily approachable, providing instant results for every possible user; a slightly less convenient one would require some basic knowledge (i.e., reading the user manual). An even more inconvenient one will require specific manual skills.

The “control” in this context will be used as a measure of how well the end result matches the user’s expectations.

The concept of control–convenience spectrum can therefore be defined quite simply in the following form:

For any creative practice that utilizes tools, the ease of operation is inversely proportional to the level of control over the creative process.

Incidentally, the word “spectrum” implies there is a continuity between the two extremes (full control and maximal convenience) and a choice of a specific tool or creative strategy determines the proportion between those extremes for any given situation.

I originally formulated the concept in early 2023³ and—at least per my limited and possibly biased outlook—it still remains relevant, besides the extreme pace at which the AI-powered solutions have advanced since then. I believe the control–convenience spectrum might itself be considered a tool for finding a right balance when working with technology-powered systems such as generative AI.

However, in order to discuss AI tools and their usability for creative practitioners, as well as describe the control–convenience spectrum concept in detail, we need to agree upon—at least temporarily, for the sake of this discussion—clear definitions of another couple of terms: “artificial intelligence” and “tool”.

3. What Is Artificial Intelligence (Or What Does AI Mean Today as Opposed to the Past)?

“Everyone is talking about AI now” might be an overstatement, but only slightly so. After public access to Chat GPT 3.5 was made available (Walsh 2025), AI became a household term basically overnight. This might be one of the reasons why the general public tends to associate the term with the latest wave of generative systems, based on artificial neural networks, or, more specifically, Large Language Models (LLMs). In the common perception, the technology is novel, dating back to late 2022 at most, even though the term itself leaked into popular culture well before the current AI bloom (such as in the case of Steven Spielberg’s movie from 2001, simply called “A.I. Artificial Intelligence”). But what exactly is this “AI” species everyone has been talking about ever since?

The term “Artificial Intelligence” had different meanings in different time periods, reflecting the fact that certain technological approaches that were once considered AI aren’t regarded as such anymore. As indicated by Sarah Gnoth and Jasminko Novak,

“The fundamental complexity of AI represents a significant challenge for public understanding. Even experts struggle to reach consensus on a unified definition of AI [...] The opaque nature of AI systems, often termed »black boxes,« further complicates users’ understanding of how AI works”. (Gnoth and Novak 2025)

A similar observation was made by Fatemeh Alizadeh et al.:

“(...) defining AI is confusing, not only for the users with no computational background, but even for the researchers and AI practitioners. This is due, among

other factors, to the evolution of the term over time and to the fact that it has never represented a single specific technology in a single specific time period”. (Alizadeh et al. 2021)

Among other factors, a so-called “AI effect” is at play here. It is the general tendency to regard technological solutions—especially those that are yet to be achieved—as “intelligent”, as long as we cannot understand how they work, and dismiss them as such the moment we start to understand the mechanics behind them. Rodney Brooks described the phenomenon in the following way:

“Every time we figure out a piece of it, it stops being magical; we say, »Oh, that’s just a computation.«”. (Kahn 2002)

John McCarthy, who was among the people who coined the very term AI in 1956, described the phenomenon in a very concise form:

“As soon as it works, no one calls it AI anymore”. (Vardi 2012)

The simplest example of this phenomenon would be an electronic calculator—while at first encounter it might seem “magical” and smarter than humans, we soon realize it just follows a simple algorithm. At the same time, in a very narrow field of basic mathematical operation, it greatly surpasses human abilities, thus it is in fact superintelligent, although to a very limited extent.

The famed Deep Blue supercomputer is another example—at first it seemed shocking that a computer could beat a chess master, but when it became widely known that the software used the brute-force approach, essentially relying on an extensive library of every possible chess move and an enormous computing power, the conclusion could only have been “that’s not real thinking” (Vardi 2012).

This shifting perception of what AI is and is not might be a reason behind the current association of the term with the newest wave of Generative AI and Large Language Models (LLMs), while at the same time disregarding any less sophisticated technologies that were once considered “intelligent”. However, the significance of the currently dominant artificial neural network approach to AI was in fact marginal throughout most of the AI research history. At the same time, approaches that today are not necessarily considered by the general public as true AI (i.e., reasoning machine), such as the symbolic approach, and, for a brief period, behavior-based robotics, dominated the field. (Garnelo and Shanahan 2019; Wahde 2016).

Let us briefly summarize the concepts behind these three approaches to AI (although there is a number of other approaches as well).

Symbolic AI—the “classic approach”—relies on a human-entered database of facts and a set of logic operations to manipulate that data. As Marta Garnelo et al. note:

“A symbolic AI system works by carrying out a series of logic-like reasoning steps over language-like representations. The representations are typically propositional in character (...) An important limitation of symbolic AI relates to (...) the extent to which its representational elements are hand-crafted rather than learned from data”. (Garnelo and Shanahan 2019)

Even though symbolic AI is based on fixed datasets and cannot spontaneously learn new facts, it is actually well-suited for generating musical structures (in terms of musical notation or MIDI data) as the tonal music, based off a relatively strict set of rules, can be easily formalized or expressed in mathematical language. A seminal example of an early AI-generated musical piece is the Illiac Suite from 1957 (Funk 2018). Even much more recent auto-accompaniment devices and software, such as “arranger keyboards” (auto-accompaniment synthesizers) and Band-In-A-Box software (Wilkins 2017) were based

on fixed algorithms (often called “styles”), so they could be considered an extension of the early symbolic-AI based music generation tools.

Therefore, the fact that computers can autonomously create music is not necessarily novel; however, the quality and character of generated musical content (complete recording in a form of a digital sound file rather than rudimentary musical notation) made the results of recent AI-based music generators penetrate the mainstream, whereas the earlier experiments went mostly unnoticed by the general public, apart from occasional mentions in the media. Most importantly from the “AI effect” perspective, tools such as the symbolic music generator used for “Iliac Suite” are not considered AI anymore. We know all too well “it is just computation”.

As a side note, symbolic AI representation of musical language requires fine-tuning of many parameters until it produces satisfactory musical results; in this regard, it bears similarities to generative art (to which I will refer later in the paper).

Another approach—Nouvelle AI, or behavior-based robotics—is based on the assumption that, rather than constructing a symbolic model of the world, a device (e.g., robot) might use various sensors to probe its actual surroundings and react to the changing circumstances based on a set of simple rules. This approach was briefly a dominant field of research in the 1980s, as Rodney Brooks was the head of Computer Science and Artificial Intelligence Laboratory at MIT (Dennis 2024) and, one of the earliest widespread implementations of AI-operated robots—the famed Roomba vacuum cleaner—was in fact based on this approach. It therefore might seem fair to consider Brooks entitled to his skepticism towards current research on artificial neural networks-powered humanoid robots (Miller 2024). In fact, even much earlier devices, such as robotic interactive installations *Senster* (Zivanovic 2005) and *Sam* by Edward Ihnatowicz, displayed seemingly autonomous behavior, based on similarly simple principles, albeit using analog techniques, rather than digital, algorithm-based systems.

The Artificial Neural Network approach is, in turn, inspired by a human brain structure; the basic assumption is that a network of artificial neurons, with the ability to rearrange the connections between them, should be able to “learn” on its own, when provided with a sufficient amount of data. In actuality, to achieve the efficiency of current artificial neural networks, a number of additional techniques had to be incorporated, such as back propagation or human-reinforced learning. These techniques obviously are not modeled on any processes observed in real, biological neural networks. When comparing the neural network approach to the symbolic approach, Garnelo and Shanahan observe that “By contrast, one of the strengths of deep learning is its ability to discover features in high-dimensional data with little or no human intervention” (Garnelo and Shanahan 2019).

Even though for many years this approach did not yield notable results, a significant step towards the bloom and eventual domination of the artificial neural network approach was the seminal paper by Rumelhart, Hinton, and Williams (Rumelhart et al. 1986) describing the backpropagation method, which led to significant improvement of this AI technique. This, together with availability of sufficient computing power, eventually led to the current AI bloom.

That was, however, a process that took another couple of decades and some significant achievements went largely unnoticed—for instance, Google Translate switched to neural-networks-based technology as soon as 2016 (Stahlberg 2020). Similarly, audio restoration software *iZotope RX* started using machine learning techniques as early as in 2017, in the *RX 6 Advanced* version (Wichern 2017). AI was also utilized by creative practitioners way before 2023, e.g., the main theme of the 2017 edition of *Ars Electronica Festival* was AI⁴. However, all these events happened before the current “media hype”, so they were not widely publicized.

Therefore, it is important to note that the current wave of generative AI systems—while significantly more powerful and developing at a much faster pace—is not in fact a paradigm shift that happened overnight, but a result of a long evolution, and a different perception is mostly a result of a sudden shift in media prominence of the topic of AI. Gnoth and Novak write about this phenomenon, citing the “missing awareness” among the common misconceptions in the “folk perception” of AI:

“(. . .) users frequently lack awareness of AI’s presence in their daily applications. This missing awareness is likely linked to common misconceptions about AI, hindering recognition of AI in everyday platforms and devices”. (Gnoth and Novak 2025)

The main downsides of the Artificial Neural Network approach are the problem of hallucination (Xu et al. 2024) and the problem of interpretability/explainability. Some other approaches to AI—i.e., symbolic approach, as exemplified in the CYC project (Sharma and Goolsbey 2019)—avoid both of these downsides, at the price of lower efficiency and flexibility.

The interpretability problem is somewhat related to the “black box” term, which refers to a device or piece of software, whose internal mechanisms are not fully understood by the user, but the user has a sufficient general grasp of its basic means of operation (i.e., which type of input will result in a specific output) (Haskel-Ittah 2023). The complexity of modern artificial neural networks makes it close to impossible to penetrate the route from a given input to an output in a reasonable timeframe (Castelvecchi 2016). In this regard AI systems are black boxes, but the important difference here is that they remain as such even for their own creators. This is also true for other, non-AI complex systems (Ladyman et al. 2013); this particular quality makes such systems distinct from the majority of other human inventions (although there were points in history when humans learned how to exploit certain phenomena—e.g., fire—long before they managed to understand them). This begs an almost philosophical question: are artificial neural networks in fact invented, or discovered (Lem 1973)?

It is important to note that the “AI effect” does not seem to apply universally to the current wave of AI—even though there is a growing understanding of how artificial neural networks work, it still tends to be viewed as “intelligent” in the common sense of the word (Gnoth and Novak 2025). This might be due to two facts:

- the level of sophistication of content created by current AI systems makes it—in certain cases at least—indistinguishable from content created by humans;
- the complexity and size of the most advanced neural networks make it impossible to penetrate the actual operation even for their creators or—generally—experts in the field.

Of course, that does not signify universal acclaim; the current generation of AI is already being dismissed as merely a “statistical machine” or “stochastic parrot” (Bender et al. 2021), therefore not a “thinking machine”. Furthermore, some experts (namely Yann LeCun) point out that, while current LLMs are doing a decent job at simulating human-like thinking, humans use the language to describe abstract thoughts and ideas (which form a basis of thinking), whereas LLMs are purely language-based, with no higher-level abstract ideas or thoughts involved (LeCun 2025).

4. AI vs. Generative Art

Generative AI should not be confused with Generative Art, which is a much older and broader term, describing specific artistic practices that utilize some kind of autonomous or automated process. Amongst many definitions, possibly the most often quoted (Galanter 2016), and a rather extensive one, was formulated by Philip Galanter:

“Generative art refers to any art practice where the artist uses a system, such as a set of natural language rules, a computer program, a machine, or other procedural invention, which is set into motion with some degree of autonomy contributing to or resulting in a completed work of art”. (Galanter 2003)

On the other hand, there is a concise definition by Sol Lewitt:

“The idea becomes a machine that makes the art”. (LeWitt 1967)

While reflecting on my own artistic practice over the last couple of years, I have come up with a short description that might as well serve as my own personal definition of generative art:

“Creating conditions for the music to emerge organically by making specific tools, rather than composing fixed pieces of music”. (Cybulski 2024)

All these definitions imply that the artist, even though they do not control all the elements of the emerging artwork directly, is indeed the creator of the artwork by designing, initiating, and controlling the generative process itself. This is an important distinction that will come in handy as we discuss the relations between artists and their tools in regard to AI.

In response to the results provided by the generative process, the artist can modify the details of the process until satisfactory artistic results are achieved. This iterative process of creating the artwork in multiple steps, a constant dialogue with the tool and matter, is by no means limited to generative art (although the boundaries between generative and non-generative art are fuzzy and may depend on subjective interpretation, as we will see later, while discussing the control–convenience spectrum in the context of visual arts).

The benefit of the generative approach is the ability to create artifacts with a high degree of complexity without the need to directly oversee all of their individual components, while still maintaining the control of the general direction of the artwork.

Last but not least, as Galanter notes, referring to the complexity theory, “In terms of our human ability to extract meaning from a given experience we require a mix of surprise and redundancy, i.e., a signal somewhere between extreme order and disorder” (Galanter 2016). Galanter calls this optimal combination of predictability and unpredictability “effective complexity”. This feature is a crucial element of a satisfying work of art. In this regard, a key part in tailoring an aesthetically successful and “effectively complex” generative system is finding a perfect balance between the elements that provide order and introduce chance.

In this regard, an AI-based tool could be considered a subsection of a generative system, as long as it is consciously used. Arguably, with the earlier generative AI systems, such as DALL·E mini, the more the system malfunctioned (i.e., hallucinated), the more artistically interesting the outcome was. This can be viewed through the lens of the effective complexity concept—the hallucination can be treated as an element of chance introduced into a broader generative system.

5. The Mainstream Perception of AI as a Music Tool

There is a discrepancy between the marketing claims of generative AI providers and the actuality of creative workflows involving digital tools, including those that could be marked as AI-powered. Arguably the most-well known generative AI system for music, Suno, already has a long back catalogue of such claims. Suno’s own social media channel “waveform.watch” regularly posts exaggerated statements, such as “Trend alert: Suno becomes go-to music tool for creators on Tiktok and Youtube” or “The Internet’s new music tool? Suno just made everyone a songwriter”. The music producer Timbaland is quoted to say about Suno: “I haven’t been excited about a tool in a long time [...] It’s the way you’re going to create music [...] It’s the new age of music creation and producing” (Mullen 2024).

The main issue here is that the general public is led to believe by such claims that anyone can be an artist/musician now, provided they use a certain AI “tool”, whereas the truth is much more complex.

Granted, using digital tools almost always implies that at least a part of the workflow is considered a black box. This is also true for other tools such as acoustic instruments—one does not need to understand all the complex mechanics of, say, a violin (such as the interplay between the top and bottom plate, bridge, soundpost and the air mass enclosed in the instrument, which effectively form a coupled system—in fact very difficult to analyze even for experts) (Chaigne and Kergomard 2016) in order to effectively use it as a tool for musical expression. Nevertheless, in order to maintain control over a tool, there has to be a delicate balance between what is understood and consciously controlled, and what is taken for granted as an intrinsic quality of a given tool.

Gnoth and Novak call this phenomenon “Mystification of AI”:

“The public perception of AI has been considerably shaped by its portrayal in media and entertainment, where it was often depicted as robots with superhuman abilities. This sensational representation attracted the excitement of many, but it may have distorted public understanding of AI’s actual capabilities, leading to an overestimation known as the “superhuman human fallacy” (...). Although generative AI systems like ChatGPT have made recent advancements in AI more accessible, they do not aid users in understanding the foundations of their decisions and outputs, complicating non-expert users’ ability to form accurate mental models of such AI systems”. (Gnoth and Novak 2025)

Apart from the alleged novelty of the relations between AI and music, some statements that are targeted towards the creative practitioners are implying that—in the face of the current revolution—they have to adapt their workflows to accommodate AI tools in order to stay competitive and not get left behind. This notion occasionally penetrates even the academic circles; for instance, it was present in at least one lecture at a recent conference on music and AI at the Chopin University of Music⁵.

6. What Is a Tool?

Generative AI systems, such as Suno, are often described as “AI tools for music”, but even a quick survey of definitions of the “tool” term begs a question: is such description justified?

Oxford Languages Dictionary defines tool as:

“A device or implement, especially one held in the hand, used to carry out a particular function.”

Tool can be also defined as:

“The use of physical objects other than the animal’s own body or appendages as a means to extend the physical influence realized by the animal”. (Jones and Kamil 1973)

“An object that has been modified to fit a purpose ... [or] An inanimate object that one uses or modifies in some way to cause a change in the environment, thereby facilitating one’s achievement of a target goal”. (Nussbaum 2011)

“A tool is an object that can extend an individual’s ability to modify features of the surrounding environment or help them accomplish a particular task [...]”. (Wikipedia Contributors 2025)

More specifically, a specialized type of a tool, i.e., a musical instrument, can be also described as a natural extension of the body (Nijs et al. 2013). Distilling it down to a simplest possible description, we may define “tool” as:

An external object, handled and controlled by a human operator to extend their natural abilities.

Handling and controlling a tool require a certain level of understanding of how the tool works, and internalized physical ability (obtained through trial and error/practice, leading to a certain degree of experience), but also a conscious intent of what the tool would be used for in the first place. This in turn stems from a general understanding of the context in which the tool is being used—i.e., chopping down the tree with an axe (how is it supposed to be done? Where does one stand in order to not be crushed by a falling tree? Is that specific tree appropriate for my needs? What exactly are my needs regarding the tree in general?); placing paint on a support with a brush (what is it that I am going to paint? What type of support am I going to use—canvas or maybe a wooden panel? What kind of paint am I using? Where exactly do I place specific colors to reproduce my visual concept? Most importantly—what is the concept?); amplifying the lip buzz with a trumpet, etc.

Having said that, a perfectly obedient tool is not necessarily beneficial to the creative process. As already proposed in the previous section, creating an artwork in a dialogue with the tool is prevalent in, but not limited to, generative art. As observed by Brian Eno, there are two distinct attitudes towards the role of the tool in the creative process; their representatives could be described as “architects” and “gardeners” (Eno 2011). In my PhD thesis I have proposed a similar distinction between a hypothetical “Composer A” and “Composer B”: the former uses an obedient tool to implement their precise concept, while the latter engages in a dialogue with the tool and matter (Cybulski 2024). Katya Davisson describes similar categories as “narrow-sense” and “broad-sense” composition (Davisson 2022). Don Norman is also a proponent of iterative process for design as a dialogue between the designer’s objectives and the external world (Norman 2013).

7. Back to the Control–Convenience Spectrum

One of the most pronounced examples from my own artistic practice of moving back and forth on the control–convenience spectrum (or a C–C spectrum) is the OSC controller called Dodecaudion (Figure 2)—a first project of panGenerator from 2011. We have created Dodecaudion as our reaction to the first wave of laptop music (laptronica) which was very prominent in Poland at the turn of the first two decades of the XXI century. Using the laptop as a musical instrument was a novel approach at that time, but the laptops were often used “saute” or with very basic MIDI controllers of the time (using mostly knobs and faders).



Figure 2. The Dodecaudion—a perfect “blank canvas” for experimentation on various sensor-to-sound mappings.

Dodecaudion sought to provide a platform for expressive musical performance, which—through the design of the device itself—would be able to fulfill the conditions of visual–auditory correlation (Jensenius 2022; Dunning 2024).

In actuality, Dodecaudion was a blank canvas—merely a controller, devoid of any built-in sounds, gesture-to-sound mappings, or performance scenarios. To become an instrument, it had to be supplemented with musical substance—both in terms of how it was supposed to sound, but—more importantly in this regard—how the proximity sensors embedded in the facets of the controller were to be mapped to those sounds. In fact, deciding upon the movement-to-sound mapping could be seen in this case as an equivalent of interaction design. In practice it was done by creating custom patches in Native Instruments Reaktor and Pure Data dataflow programming environments.

Let us take a closer look at some of the most distinguished sets of those mappings:

- Case 1: Direct mapping of each facet/sensor to a specific note from a single octave of the 12-tone equal tempered scale. (video example: <https://youtu.be/cbkOvw9Tv04>, accessed on 15 December 2025).

As the name implies, Dodecaudion is a regular polyhedron with 12 faces, each containing a single proximity sensor. Therefore, this mapping seemed a natural starting point for the exploration of musical possibilities of the instrument.

The 12 faces would then reflect a single octave of a piano keyboard, although the velocity of each note could be controlled continuously, not unlike a mixing desk with a set of 12 faders controlling the volume of the notes from a chromatic scale. In reality, this mapping proved difficult to master; even physical faders would be less convenient than a regular piano keyboard. The faders, however, would at least allow a certain combination of notes to be set up and sustained. The sensors on the Dodecaudion do not facilitate the “hold” functionality; in order to sustain a note of a given volume, the performer has to hold their hand steadily in a certain proximity to the sensor. Moreover, there is no physical contact with the instrument; it imposes difficulties that are always associated with this kind of “air” instruments, as observed already in 2000 by Rován and Hayward (2000). The lack of tactile feedback is a major downside of this species of instruments, rendering them difficult to master. Additional nuisance stems from the fact that, while reaching out to one sensor, the performer may accidentally trigger another one, thus “hitting the wrong note”. Performative difficulties apart, this mapping also did not yield very interesting musical results.

- Case 2: panGenerator was asked to make Dodecaudion available for the El Hormiguero talk show on Spanish TV. I have prepared a fresh set of sounds and mappings, as the instrument was to be played by inexperienced performers and the entry barrier had to be as low as possible. (video example: <https://youtu.be/Sco8d4xpWR0>, accessed on 15 December 2025).

In this case, the data from all sensors was averaged to a single parameter, which gradually changed the volume and filter cutoff on three virtual instruments—synthetic drums, bass, and harmonic pad. Other parameters of the instruments were in turn controlled by a sequencer, preprogrammed with a looped four-chord harmony and simple beat.

This mapping forms a polar opposite of the one used in Case 1: it facilitates playing music instantly, despite a lack of expertise either with the Dodecaudion itself, or with musical instruments altogether. However, it soon becomes clear—both to the performer and the audience—that the performer is not really playing the instrument in terms of controlling the actual musical parameters, such as pitch, tempo, etc., but rather just controls the mix and timbre in a very coarse fashion. Therefore, even though

seemingly compelling musical results can be obtained effortlessly, it has very little in common with playing an actual musical instrument.

- Case 3: At first glance, similar to Case 1—each sensor controls a single sound. In reality, the mapping is slightly different: the distance from the sensor influences the note pitch, quantized to an arbitrary musical scale (each of the sensors produces slightly different pitch range), rather than controlling the volume, which in this case is influenced by the speed of hand movements—the faster, the louder. Thus, more expressive movement correlates clearly with the volume and the rate of note changes. As the notes triggered by all the sensors come from the same scale, there is no risk of playing “wrong notes” when accidentally hitting an adjacent sensor. (video example: <https://youtu.be/yZsJjVSToCk>, accessed on 15 December 2025).
- Case 4: In order to create more complex, but still controllable musical structures, I have created a simple semi-autonomous system, consisting of a generative beat machine, a Karplus–Strong type synthesizer with sounds triggered by the drum samples, and an “8-bit” style synth, controlled by a generative sequencer. A couple of parameters (such as complexity of the beat, pitch of the drum samples, randomization of the sample choice, Karplus–Strong “string” decay and high frequency damping) were assigned to only six out of twelve sensors (the rest of the sensors remained inactive to avoid accidental triggering). In such an arrangement, all the parameters could be controlled at once. Additionally, some parameters were assigned to presets, switchable by an external foot controller. (video example: https://youtu.be/_Bn_pEKEDZM, accessed on 15 December 2025).

Such mapping provided enough control for the Dodecaudion to really become an instrument—i.e., providing the performer with enough influence on the music so they can feel in charge, as well as creating convincing sound–action coupling for the audience (Jensenius 2022), but enough convenience to prevent the performance from becoming a tedious ordeal, as in the Case 1.

Working on Dodecaudion provided me with a unique opportunity to practically examine how a single hardware controller can be perceived in significantly different ways, both by the interactor (performer) and the audience, depending on the type of gesture-to-sound mapping applied. An interesting observation emerges already from this first example: the most interesting musical outcomes, as well as the most satisfying user experience for the performer, stem from the third and fourth cases—a mapping that is somewhere in a narrow sweet spot between the two extremes—“inconvenient control” and “lazy convenience”. Notably, Case 2 is obviously rather extreme, exemplifying how the goal of creating an instrument that would be appealing to an untrained performer on one hand, and mass audience on the other, fails to produce both meaningful interaction and compelling aesthetic results. This is, however, precisely the promise of many of the commercial generative AI systems—that, regardless of the level of experience, with the help of a given system anyone can produce captivating art.

We may now place the four mappings of Dodecaudion on the control–convenience scale in such a way (Figure 3):




Figure 3. Various Dodecaudion mappings on a control–convenience scale.

In a similar fashion, we could arrange a number of different, but similar instruments by their control/convenience ratio (Figure 4):

- Cobza is a traditional Romanian plucked chordophone ([Hornjatkevych 2021](#)), with a multistring arrangement, which is already a technical improvement over simpler, single-stringed instruments, as it extends the range of the instrument. A fretless neck allows for an unlimited choice of musical scales (i.e., microtonal, xenharmonic, etc.), also facilitating glissandi and portamento; however, it requires more skill than fretted instruments. Playing chords in tune also poses some challenges.
- The introduction of frets on instruments such as the guitar takes away a lot of the effort of playing in tune, “quantizing” the pitch to a predesigned set of chromatic notes, offering more convenience, but also taking away the ability to play alternate tunings, glissandi, etc.
- Strumstick ([McNally 2025](#)) goes one step further towards the ease of playing, narrowing down the note choice to a diatonic scale by a specific fret placement and tuning; at the same time, it takes away the ability to play in different keys than those planned by the designers.
- Autoharp ([Kettlewell and Long 2013](#))—in all previous instruments, the left hand of the player was in direct contact with the strings, which allowed to control the timbre and sustain of the notes, as well as produce vibrato and glissandi. However, playing chords required training to achieve muscle memory and technical ability. Autoharp removes that necessity by introducing a set of levers with felt mutes, arranged in predefined chord configurations, thus facilitating chord change by pressing the levers. The tradeoff is the inability to control the string vibration and pitch directly, which was an inherent quality of all the previous instruments.




Figure 4. Various string instruments on a control–convenience scale.

We could also see how an evolution of a single instrument makes it move through the C–C spectrum: in case of the trumpet, a simplest form—a natural trumpet—was already a technological achievement, which helped to amplify the lip buzz, and—by the means of the resonant frequencies of the air column within the trumpet pipe, made it possible to play discrete notes within the natural major scale (stemming from the harmonic series). Introduction of a slider facilitated playing the notes from outside the harmonic series but required good kinesthetic coordination. Quick note jumps, possible to achieve with lip-trills used in the “clarino” technique ([Dahlqvist and Tarr 2001](#)), are however not possible to the same extent on a slide trumpet. The introduction of the valves moved the instrument further towards the convenience end; however, it removed the ability for playing glissandi, yet again sacrificing some control for more convenience.

For the last example in this section, let me look at the visual arts through the lens of the C–C spectrum concept (Figure 5). Galanter argues that Jackson Pollock’s technique should not be considered generative art, since it relies on the properties of the physical world (i.e., dripping paint) and as such, all art should be considered generative ([Galanter 2003](#)). There is, however, a varying degree to which different painting techniques rely on the physical properties of the paint: the traditional oil painting technique strives for full control over the medium. Watercolor painting, in turn, relies heavily on the diffusion of pigments in water, utilizing the process to create nature-like effects. This, in turn, is not so different from Max

Ernst's technique of decalcomania, which might be considered a full-fledged generative art method. We could then attempt to create another graph, imposing various painting techniques against the control–convenience axis.

This example, however, is not as straightforward: watercolor painting might be a step towards convenience by the means of its semi-generative characteristics. It is nevertheless a difficult technique to master. So, on the following graph, we will focus on the ease of creating life-like images with a given technique, rather than the ease of handling the actual medium. In this context, the most convenient method of producing such images would be using the “paint-by-numbers” kit (Langer 2019):




Figure 5. Various painting techniques on a control–convenience scale.

As we could already see from the last couple of examples, the C–C spectrum concept obviously has its limitations. For instance, an instrument designed for the ease of creating simple melodies or chords (i.e., Strumstick or Autoharp) might be very convenient as an entry-level instrument for an unskilled musician, while at the same time posing huge obstacles for a skilled musician keen to play specific scales or chords, which are not available on these instruments. At the same time, the latter might feel completely at ease playing desired harmonies on a “less convenient” instrument, such as a guitar or a cobza.

For yet another example of a seeming failure of the C–C spectrum, we might look at the violin—an instrument notorious for its steep learning curve, as producing even a single sonically acceptable note (at least in the common perception of “nice” and “unpleasant” timbre) requires considerable skill. Thus, the user interface of the violin is an inconvenient one for an unexperienced user. Nevertheless, a skilled violinist approaches the instrument with ease and producing a fine sound does not require conscious effort anymore. This would suggest that the same instrument offers different mixture of control and convenience, depending on the skill of the user. Indeed, eventually even a most demanding instrument becomes more approachable, to the point that playing it becomes easy. Nevertheless, the effort and time invested in achieving proficiency on a given instrument might be viewed as the “inconvenient part” of the whole musician–instrument interaction—thus, in order to gain more control, one has to commit themselves to the process by acquiring the skill necessary to operate the tool at a desired level of control. This notion will be also touched upon in the context of commercial AI systems in the last paragraph of the Section 8.

8. Control–Convenience Spectrum and Generative AI

As I have already proposed in the “What is a Tool?” subsection, handling and controlling a tool requires a certain level of understanding of how the tool works, internalized physical ability to operate it, but also a conscious intent of what the tool would be used for. This leads us to two major realizations regarding the status of generative AI systems:

- in order for it to be considered a tool, it would have to allow a certain amount of control over the final result, but:
- in order to control a tool, one has to possess both some experience in handling the tool and a wider knowledge of the discipline/task, for which the tool is being used.

In this regard, the claim that “generative AI tools will make everyone an artist” is deeply flawed. One could argue that AI systems such as Suno (or Suno Studio/Suno Pro for that matter) allow for a deeper degree of control, becoming a real tool in the process.

But moving the control–convenience slider towards the “control” end moves it away from the “convenience” end by requiring some amount of skill and understanding of the musical matter. Thus, there is no “magic button”; if we want to obtain what we want from generative AI, we have to know what it is that we want in the first place, and the more complex the task at hand, the more precise and informed the prompt/instruction must be. It is exactly the same as with any other tool, and AI tools are no different in this regard.

In a real-life example, a friend of mine wanted to add mechatronic elements to their acoustic instrument; having only a vague idea on how it is supposed to work, they asked Chat GPT for assistance. Unsurprisingly, the answers produced by the LLM were as vague as the prompt itself, and the details of specific technical means varied from reply to reply. For instance, Chat GPT suggested that an electric motor requires a separate power source, but the actual parameters (i.e., voltage) changed with each reply. It might as well have asked for more details of the motor, and only then tried to match the power supply to that specific device. Needless to say, the lack of consistency in the replies provided by the LLM only added to the initial confusion of my friend.

This shows us once again that—in order to create more complex projects in cooperation with LLMs or Generative AI systems—one needs to possess certain knowledge and expertise in the field, in order to formulate the prompt in a detailed fashion. Even then, though, the replies provided by AI have to be scrutinized for possible errors and hallucinations, which are inherent features of all artificial neural network-based systems.

Nevertheless, there are areas where LLMs can produce sensible results even from very vague prompts. Chat GPT 5 provides decent, or at times even very good code, as long as it is to be written in one of the more popular programming languages. For instance, a widespread availability of JavaScript code examples on the web is most likely the secret behind Chat GPT’s proficiency in this particular language (i.e., there was enough free training data for the LLM to become proficient). In this area of expertise (unlike the example with a power supply), Chat GPT can ask the right specifying questions, should the prompt be too vague or illogical. It can also produce mid-code comments, so various sections of the LLM-produced code can be easier to understand and modify by the prompter.

Obviously, modifying and tailoring the code to real life applications requires a certain grasp of coding knowledge from the user. Even a simple realization “OK, what you provided me with is fine, but that is not exactly what I wanted” is an attempt to gain more control, and at the same time is already a step away from the “convenience” end of the spectrum. If the prompt is not precise enough, the LLM might ask additional questions or provide the user with options, from which they could choose, eventually supplying LLM with a more precisely formulated task—all as a result of the guiding questions from the LLM.

One might conclude, then, that my earlier claims are therefore refuted—you actually can start off with no skills or understanding whatsoever, and the AI tool will lead you through the process of refining your initial prompt. The truth is, however, that through the process of considering the questions, options, and suggestions from the AI tool, the user actually gains knowledge and experience, which is eventually utilized to create the final, refined prompt. It also takes more time to get into discussion with the LLM and consider the options provided. Therefore, the user does not move towards the “control” end of the spectrum “for free” (i.e., without sacrificing at least some convenience). We cannot “have a cookie and eat a cookie”.

9. But Why Would We Prefer Convenience over Control?

9.1. Do Artists Really Desire the Creative Process to Be Easy?

In early 2025, Mikey Shulman, the CEO of Suno, had been quoted to say the now-infamous words:

“It’s not really enjoyable to make music now. . . it takes a lot of time, it takes a lot of practice, you have to get really good at an instrument or really good at a piece of production software. I think the majority of people don’t enjoy the majority of time they spend making music”. (Patti 2025)

This statement was met with significant media backlash. For instance, Luis Prada remarks:

“Besides coming off as deeply out of touch with the spirit of creating music, Shulman comes off as a guy doing a piss-poor job of selling his product. He invented a problem that no one has and is trying to position his product as the solution”. (Prada 2025)

Also noticing another significant issue:

“The irony here, of course, is that Suno AI runs on generative AI tools trained on copyrighted music created by people who actually are talented, people who did not provide consent for their work to be used to train AI models and are now suing Suno because of it”. (Prada 2025)

As Daniel Griffiths rightfully observes, Shulman implies that “(. . .) music production is a chore that Suno can rid us of the need to do” (Griffiths 2025), and counters that implication by stating that:

“Music making is a process, a task that needs to be accomplished, a triumph over adversity, and his software can help. But, much like riding a motorbike in a marathon or having a pro chef prep your dinner on *Come Dine With Me*, could this be assistance that rather ruins the point of taking part at all?”. (Griffiths 2025)

Eve Upton-Clark picks up on an interesting notion in the Schulman–Patti interview, for instance:

“(. . .) the interviewer attempts to compare the creative process to running: While it might not be enjoyable at first, once the necessary muscles are built up, people often fall in love with it. Shulman counters that most people stop running before this can happen. Rather than pouring time and effort into practicing and honing their craft, generative AI platforms like Suno allow anyone to produce full songs from just a few prompts, bypassing the traditional barriers for entry entirely”. (Upton-Clark 2025)

The article also quotes an anonymous X user, saying that people like Shulman:

“(. . .) have always been envious of those with commitment and talent and love that now they don’t need to ‘waste time’ perfecting and nourishing a craft. They miss the whole point and just want a ‘I win’ button in life”. (Upton-Clark 2025)

This last reaction especially resonates with the concept of flow, originally formulated by Mihály Csíkszentmihályi as “(. . .) a state of concentration so focused that it amounts to absolute absorption in an activity” (Csíkszentmihályi 1990). Flow is accompanied by “(. . .) a sense of deep enjoyment that is so rewarding people feel that expending a great deal of energy is worthwhile simply to be able to feel it” (Csíkszentmihályi 1990).

Flow tends to happen during creative activities, most likely because they fulfill basic conditions of the state of flow:

“Optimal experiences are reported to occur within sequences of activities that are goal-directed and bounded by rules—activities that require the investment of psychic energy (attention) and that could not be done without skills”. (Csíkszentmihályi 1990)

Therefore, achieving the flow state requires meeting certain conditions (Figure 6), which generally come down to balancing the level of difficulty of the task being performed with the level of skills needed to complete it (Cybulski 2024).




Figure 6. Acquiring the state of flow requires proper alignment between the challenge level and the skill level (Csíkszentmihályi 1998).

Csíkszentmihályi developed the concept in a number of publications and the description above is a far-reaching simplification of his thought, but for the sake of this discussion it provides yet another argument (apart from desire to maintain control) for why creative practitioners tend to despise the idea of a “magic button” and, instead, seek to derive satisfaction from the creative process, even though it requires certain investment of energy and skill.

9.2. Do We Really Need More Music to Be Produced Daily?

Last but not least, the notion that “everyone can be a musician now thanks to AI” is, as I have tried to convey, not true. However, whether it would be a good thing in the first place, is at the very least debatable. As Griffiths notes, “SunO makes it child’s play to produce music that—to the untrained ear at least—could pass for... music” (Griffiths 2025). Thus, even though the technology can produce convincingly sounding music, it is, at least for now, generic and unoriginal at best, meeting the criteria for “AI slop”. It is, however, already being used in less-crucial applications, which call for any sort of muzak/elevator/background music. The use for AI-generated music as a playlist filler for streaming platforms is also expanding. Therefore, even though the artists who take their business seriously should not be worried just yet, AI will most likely soon replace musicians who are lowest on the food chain, working day jobs producing generic “content”.

Arguably, the area where musicians are the least likely to be replaced by technology, is “live music performance” defined as an act of singing and playing instruments by humans. However, even though it might have been an obvious definition for the majority of the history of music, the introduction of technology (even in the form of XIX-century mechanical instruments) started to blur the boundaries between the “live” and “automated” performance. Currently, while entering a jazz club, we might still expect human musicians playing mostly acoustic instruments without the aid of technology such as drum machines or pre-recorded loops. However, such technological aids would be perfectly acceptable in a pop music context. The irreplaceability of human musicians is further questioned by performances of virtual singers (i.e., the Gorillaz or Hatsune Miku) or even entire bands

(virtual ABBA tours), gradually shifting the boundaries of what is still accepted as a live music performance.

Apart from that, a growing number of average or mediocre AI-generated music is adding up to the phenomenon of “AI slop”, elevating the noise floor of worthless content, which music consumers have to get through in order to find truly valuable artifacts. As noted by Alexandra Plesa, “With millions of new tracks dropping daily, finding something special takes more work than ever” (Plesa 2024). The article itself bears a symptomatically alarming title: “More Music Is Released Daily Now Than the Entire Year of 1989, and It’s Breaking the Industry”.

The AI-generated content has also the ability to devalue pre-existing cultural goods, such as was the case with a wave of AI-generated Studio Ghibli-style art; a facsimile of something that was revered for its aesthetically outstanding qualities for many years, became, all of a sudden, readily available at the press of a button. The value of the very art that was used as a training data for generative AI tools suddenly was found at the risk of being perceived as a passing trend.

Moreover, even though the likes of Mikey Schulman and Sam Altman encourage us to take the subject lightly and engage in innocent playtime with their generative AI products, AI-generated music is also expensive, both literally and metaphorically—unprecedented amounts of money are being invested in the development of AI platforms, but also the power needed to run the data centers consumes enormous amounts of resources. The technology also makes use of underpaid labor to execute the “human reinforced learning” method (Sixtus 2025). It might also affect the prices of seemingly unrelated products, such as consumer laptops, as some computer chip suppliers already started shutting down the consumer products lines, to focus on more profitable products tailored to the AI datacenters’ demands (Zuhair 2025). This might, in turn, create yet another backlash for creative practitioners who use laptops as their tools.

As Peter Kirn observed, the situation is not black or white, but it is definitely time for serious reconsideration of AI-powered creative practices:

“(. . .) I have interest in generative sound, algorithmic music, machine learning. It’s not about being pro- or anti-AI like this is a sport. We’re talking about the critical examination of a technology that is sucking up a huge amount of resources and reshaping the world around us”. (Kirn 2025)

10. Conclusions: Post-AI Music

Throughout the paper, I referred to several recurring themes: the control–convenience spectrum, effective complexity, and flow. These three, seemingly disparate concepts, share at least one common trait: recognizing the need for balance between extremes—control and ease, chaos and order, skill level and challenge level. Theoretical formulation of this common trait—a call for balance—is rather straightforward; achieving the equilibrium in real-world scenarios is a much more involved task, and the actual details and parameters differ from situation to situation. Even though—for a very specific scenario, say, creating satisfactory visual patterns with analog video feedback—it is possible to create precise guidelines (i.e., setting the camera zoom and iris as well as monitor brightness and contrast within a certain range) (Crutchfield 1984), the actual “sweet spot” can only be found by experimentation, direct manual involvement, and trial and error.

Thus, even though an artist working in dialogue with a tool that exhibits a certain level of autonomy may be accused of “pawning off” some of the undesired struggle to the technology, the actual creative effort lies in designing and controlling the generative process in such a way that it maintains a delicate balance between the extremes. This involves both a conscious, intellectual understanding of the inner workings of the setup, and more intuitive,

personal taste-based process of decision making—the decision in question adheres to where to set the boundaries between said extremes to achieve a satisfying aesthetic outcome that the artist considers their own personal statement.

In my PhD thesis I attempted my own definition of post-digital art, which I like to think of as an approach utilizing optimal combination of digital and non-digital elements for a given application—the flexibility of code and the tangibility of physical matter. Obviously, such a fine balance for any new technology could only be achieved after the initial captivation and enchantment with its capabilities. For the pre-AI digital solutions, this took a good couple of years, but eventually, settling on a post-digital approach became a reality for many artists. As Magnusson et al. note, “After decades of digital systems, people are increasingly wanting to explore sound and complexity in tactile form” (Magnusson et al. 2017).

The relations between creative practitioners and AI-based systems might follow the same path, as many of the issues I reflected upon in previous sections seem to indicate that we are reaching a point of exhaustion with AI-generated art. In this vein, we are most likely heading towards post-AI art—or more specifically, post-AI music—which could be defined in a similar fashion to my understanding of the earlier term, as a creative practice involving any combination of acoustic, analog, digital and AI-based techniques that is best suited for the task at hand. For most artists, this toolset will hardly ever consist solely of generative AI systems, provided by major companies.

Even though LLMs and Generative AI are attractive by the sheer size of those networks, offering seemingly unlimited possibilities, the creative process does not necessarily benefit from such circumstances, as I have tried to argue for a healthy portion of this paper. In most cases, using those tools could be compared to using a cannon to kill a mosquito. All their downsides, outlined in the previous sections, indicate in favor of less expensive, more purpose-fitting solutions.

Limited, but accurately tailored AI solutions, applied only to the tasks that could not be completed with any other means, might offer a more streamlined workflow and facilitate an iterative process of dialogue between the artist and the tool. These post-AI workflows, combining physical and analog elements with custom, streamlined neural networks trained on data provided by their creators, already exist for a good couple of years. For instance, Rebecca Fiebrink introduced her Wekinator software (v1.0) already in 2009 (Schedel et al. 2011). More recently, the research on AI and music is conducted in several research groups and universities, such as Augmented Instruments Lab⁶ and Intelligent Instruments Lab⁷, among others. The latter is conducting research on applications of explainable (interpretable) AI in music (Betancur et al. 2023), as well as AI systems that can run locally, even on small, embedded computers like Bela (Pierce et al. 2023).

Today’s technological landscape changes enormously fast; just as I was finalizing this paper, it was announced that Rodney Brooks’ iRobot company, touched upon earlier on in the paper, had filed for bankruptcy. Similarly, a remark I made on the shrinking supply of RAM chips was to include a link to a news story about chips produced by Samsung, supposedly withdrawn from the consumer electronics market by the manufacturer. Nevertheless, the news was revealed to be fake just two days before I submitted this paper for reviews. Perhaps I should then, one more time, take inspiration from Lem, who proposed in his “Wielkosc Urojona” that encyclopedias should foresee the future, in order to stay up to date as long as possible (Lem 1973).

In this vein, I will refer to the abstract of a yet to be published publication. In it, Stephen Roddy and Brian Bridges observe that “(. . .)there is a resurgence of cybernetic approaches to music-making in response to the rapid proliferation of machine learning technologies for artificial media production in recent years.” The abstract promises that the publication will include “(. . .)suggestions how an ethically grounded cybernetics of musical

machine learning might open a new space of artistic possibility beyond the algorithmic monoculture of AI slop” (Roddy and Bridges 2026).

That might indicate that the views I have expressed in this paper are not isolated and that, hopefully, we might see more artists adapt a more nuanced approach to AI. We might have just experienced a major mainstream milestone towards post-AI art, as Apple TV created their recent ident (video logo) using traditional, pre-digital, in-camera techniques (Snelling 2025). While this might be a calculated attempt at regaining reputation after a highly criticized iPad campaign from the previous year (Milmo 2024), rather than an idealistic statement, the cat is out of the bag for the first time on such a large scale. Rather than combining the best of both worlds, this is a complete retreat from generative AI and digital image manipulation altogether. At the same time, Coca-Cola produced another Christmas commercial entirely with generative AI. An article by Harry Boulton quotes Coca-Cola officials stating “(..)we need to keep moving forward and pushing the envelope (..) The genie is out of the bottle, and you’re not going to put it back in” (Boulton 2025). We have heard that one already—“you must use AI in order to stay competitive”. It is our call to decide which attitude we like better.

The future might be bright, but, while moving forward, we must take part in its creation by fostering more balanced creative practices, which will use AI as an actual tool, not a substitute for laziness or lack of creativity.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The author declares no conflicts of interest.

Notes

- ¹ Ars Electronica Archive—AI LAB, <https://archive.aec.at/ailab/>. Accessed on 15 December 2025.
- ² New Interfaces for Musica Expression conference proceedings, <https://www.nime.org/papers/>. Accessed on 15 December 2025.
- ³ I’ve expressed it at the “Sound Engineering and New Media” symposium at the AMFN, and also reflected upon it at the “AI in Music” conference at UMFC, but it was never published in a written form.
- ⁴ Ars Electronica Archives/AI-related projects, <https://webarchive.ars.electronica.art/festival/2017/ai/en.html>. Accessed on 15 December 2025.
- ⁵ AI w muzyce—Konferencja UMFC | FUTUROPOLIS, https://chopin.edu.pl/szczegoly-wydarzenia/1431_ai-w-muzyce-konferencja-umfc-futuropolis. Accessed on 15 December 2025.
- ⁶ Augmented Instruments Laboratory, <https://instrumentslab.org/>. Accessed on 15 December 2025.
- ⁷ Intelligent Instruments Lab, <https://iil.is>. Accessed on 15 December 2025.

References

- Alizadeh, Fatemeh, Gunnar Stevens, and Margarita Esau. 2021. I Don’t Know, Is AI Also Used in Airbags?: An Empirical Study of Folk Concepts and People’s Expectations of Current and Future Artificial Intelligence. *i-com* 20: 3–17. [CrossRef]
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *FAccT ’21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York: ACM, pp. 610–23. [CrossRef]
- Betancur, Celeste, Jack Armitage, Victor Shepardson, and Nicola Privato. 2023. Explainable AI in Music Performance: Case Studies from Live Coding & Sound Spatialisation. In *XAI in Action: Past, Present, and Future Applications, NeurIPS 2023 Workshop, New Orleans, LA, USA, 10–16 December 2023*. La Jolla: NeurIPS.
- Boulton, Harry. 2025. Apple Give ‘Biggest Middle Finger to Coca Cola’ with Mind-Blowing New TV Ad. Uniladtech.com. Available online: <https://www.uniladtech.com/apple/apple-give-middle-finger-coca-cola-new-ad-461384-20251107> (accessed on 15 December 2025).

- Castelvecchi, Davide. 2016. Can we open the black box of AI? *Nature* 538: 7623. [CrossRef] [PubMed]
- Chaigne, Antoine, and Jean Kergomard. 2016. Coupled Systems. In *Acoustics of Musical Instruments*. Edited by Antoine Chaigne and Jean Kergomard. New York: Springer.
- Crutchfield, James P. 1984. Space-time dynamics in video feedback. *Physica D: Nonlinear Phenomena* 10: 229–45. [CrossRef]
- Csikszentmihályi, Mihaly. 1990. *Flow: The Psychology of Optimal Experience*. New York: Harper and Row.
- Csikszentmihályi, Mihaly. 1998. *Finding Flow: The Psychology of Engagement with Everyday Life*. New York: Basic Books.
- Cybulski, Krzysztof. 2024. Improvizacja wspomagana technologią—Cyfrowo sterowane instrumenty akustyczne jako element improwizowanego utworu muzycznego. Unpublished doctoral dissertation, Chopin University of Music, Warsaw.
- Dahlqvist, Reine, and Edward H. Tarr. 2001. Clarino. In *Grove Music Online*. Oxford: Oxford University Press. Available online: <https://www.oxfordmusiconline.com/grovemusic/view/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-0000005865> (accessed on 15 December 2025).
- Davison, Katya. 2022. Improvisation as a Method of Composition: Reconciling the Dichotomy. *The British Journal of Aesthetics* 62: 373–86. [CrossRef]
- Dennis, Michael Aaron. 2024. Rodney Brooks. Encyclopedia Britannica (Online). December 26. Available online: <https://www.britannica.com/biography/Rodney-Allen-Brooks> (accessed on 15 December 2025).
- Dunning, Graham. 2024. Ironing in the Creases: Developing an Idiosyncratic Electro-mechanical Musical Instrument by Reinforcing Its Faults. Paper presented at the International Conference on New Interfaces for Musical Expression, Utrecht, The Netherlands, September 2–6. [CrossRef]
- Eno, Brian. 2011. Composers as Gardeners. Edge.org. Available online: https://www.edge.org/conversation/brian_eno-composers-as-gardeners (accessed on 15 December 2025).
- Funk, Tiffany. 2018. A Musical Suite Composed by an Electronic Brain: Reexamining the Illiac Suite and the Legacy of Lejaren A. Hiller Jr. *Leonardo Music Journal* 28: 19–24. [CrossRef]
- Galanter, Philip. 2003. What Is Generative Art? Complexity Theory as a Context for Art Theory. Available online: http://philipgalanter.com/downloads/ga2003_paper.pdf (accessed on 15 December 2025).
- Galanter, Philip. 2016. Generative Art Theory. In *A Companion to Digital Art*. Edited by Christiane Paul. New York: Wiley Blackwell.
- Garnelo, Marta, and Murray Shanahan. 2019. Reconciling deep learning with symbolic artificial intelligence: Representing objects and relations. *Current Opinion in Behavioral Sciences* 29: 17–23. [CrossRef]
- Gnoth, Sarah, and Jasminko Novak. 2025. Supporting AI Literacy Through Experiential Learning: An Exploratory Study. In *Learning and Collaboration Technologies. HCII 2025. Lecture Notes in Computer Science*. Edited by Brian K. Smith and Marcela Borge. Cham: Springer, vol. 15808, pp. 233–51. [CrossRef]
- Griffiths, Daniel. 2025. “It’s Not Really Enjoyable to Make Music Now,” Says the CEO of an AI Music-Making Platform, but Don’t Worry—He’s Here to Help. Musicradar.com. Available online: <https://www.musicradar.com/music-industry/its-not-really-enjoyable-to-make-music-now-says-the-ceo-of-an-ai-music-making-platform-and-he-wants-to-help> (accessed on 15 December 2025).
- Haskel-Ittah, Michal. 2023. Explanatory black boxes and mechanistic reasoning. *Journal of Research in Science Teaching* 60: 915–33. [CrossRef]
- Hornjatkevych, Andriy. 2021. The Kobza and the Bandura: A Study in Similarities and Contrasts. *FOLKLORICA—Journal of the Slavic, East European, and Eurasian Folklore Association* 13. [CrossRef]
- Jensenius, Alexander Refsum. 2022. *Sound Actions: Conceptualizing Musical Instruments*. Cambridge, MA: MIT Press.
- Jones, Thony B., and Alan Kamil. 1973. Tool-making and tool-using in the northern blue jay. *Science, New Series* 180: 1076–78. [CrossRef] [PubMed]
- Kahn, Jennifer. 2002. It’s Alive. *Wired* 10: 30. Available online: <https://www.wired.com/2002/03/everywhere/> (accessed on 15 December 2025).
- Kettlewell, David, and Lucy M. Long. 2013. Autoharp. In *Grove Music Online*. Oxford: Oxford University Press. Available online: <https://www.oxfordmusiconline.com/grovemusic/view/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-1002248189> (accessed on 15 December 2025).
- Kirn, Peter. 2025. Benn Jordan’s AI Poison Pill and the Weird World of Adversarial Noise. *Create Digital Music*. Available online: <https://cdm.link/benn-jordan-ai-poison-pill/> (accessed on 15 December 2025).
- Ladyman, James, James Lambert, and Karoline Wiesner. 2013. What is a complex system? *European Journal for Philosophy of Science* 3: 33–67. [CrossRef]
- Langer, Emily. 2019. Dan Robbins, Paint-by-Number Inventor Who Made ‘Every Man a Rembrandt,’ Dies at 93. *The Washington Post*. Available online: https://www.washingtonpost.com/local/obituaries/dan-robbins-paint-by-number-inventor-who-made-every-man-a-rembrandt-dies-at-93/2019/04/05/6f80f412-574e-11e9-8ef3-fbd41a2ce4d5_story.html (accessed on 15 December 2025).
- LeCun, Yann. 2025. *Geometry of Machine Learning Special Lecture, Department of Mathematics*. Cambridge, MA: Harvard University. Available online: <https://youtu.be/yUmDRxV0krg?si=sDrvs-d4WTSckF7> (accessed on 15 December 2025).
- Lem, Stanislaw. 1973. *Wielkosc Urojona*, 1st ed. Warsaw: Czytelnik, p. 122.

- LeWitt, Sol. 1967. Paragraphs on Conceptual Art. *Artforum* 5: 79–83.
- Magnusson, Thor, Chris Kiefer, and Halldor Ulfarsson. 2017. Reflexions upon Feedback. In *Proceedings of International Conference on New Interfaces for Musical Expression, Copenhagen, Denmark, 15–18 May 2017*. Geneva: Zenodo. [CrossRef]
- McNally, Bob. 2025. History of the Strumstick by Bob McNally. Available online: <https://strumstick.com/pages/history-of-the-strumstick> (accessed on 15 December 2025).
- Miller, Ron. 2024. MIT Robotics Pioneer Rodney Brooks Thinks People Are Vastly Overestimating Generative AI. *TechCrunch*. Available online: <https://techcrunch.com/2024/06/29/mit-robotics-pioneer-rodney-brooks-thinks-people-are-vastly-overestimating-generative-ai/> (accessed on 15 December 2025).
- Milmo, Dan. 2024. Apple Apologises for iPad ad Criticised as ‘Destruction of the Human Experience’. *The Guardian*. Available online: https://www.theguardian.com/technology/article/2024/may/09/apple-ipad-ad-prompts-online-backlash?CMP=share_btn_url (accessed on 15 December 2025).
- Moura, Francisco Tigre, ed. 2024. *Artificial Intelligence, Co-Creation and Creativity: The New Frontier for Innovation*, 1st ed. London: Routledge. [CrossRef]
- Mullen, Matt. 2024. “This Is Not Something to Be Afraid of—This Is Something That We Need to Use”: Timbaland Says He Spends 10 Hour a Day Using Controversial AI Music Generator Suno. *MusicRadar.com*. Available online: <https://www.musicradar.com/artists/this-is-not-something-to-be-afraid-of-this-is-something-that-we-need-to-use-timbaland-says-he-spends-10-hours-a-day-using-controversial-ai-music-generator-suno> (accessed on 15 December 2025).
- Nijs, Luc, Micheline Lesaffre, and Marc Leman. 2013. The musical instrument as a natural extension of the musician. In *Music and Its Instruments*. Sampzon: LAM-Institut Jean Le Rond d’Alembert.
- Norman, Donald A. 2013. *The Design of Everyday Things*. Cambridge, MA: The MIT Press.
- Nussbaum, Martha C. 2011. The Capabilities Approach and Animal Entitlements. In *The Oxford Handbook of Animal Ethics*, online ed. Oxford Handbooks; Edited by Tom L. Beauchamp and R. G. Frey. Oxford: Oxford Academic, pp. 228–52. [CrossRef]
- Patti, Mike. 2025. It’s Not Really Enjoyable to Make Music Now. . . , X. Available online: <https://x.com/mpatti/status/187814702976491113?s=20> (accessed on 15 December 2025).
- Pierce, Ezra, Victor Shepardson, Jack Armitage, and Thor Magnusson. 2023. Bela-IREE: An Ap-proach to Embedded Machine Learning for Real-Time Music Interaction. In *Proceedings of AIMC 2023*. Piscataway: IEEE. Available online: <https://aimc2023.pubpub.org/pub/t2l10z49> (accessed on 15 December 2025).
- Plesa, Alexandra. 2024. More Music Is Released Daily Now Than the Entire Year of 1989, and It’s Breaking the Industry. *Headphonesty.com*. Available online: <https://www.headphonesty.com/2024/11/more-music-released-daily-entire-year-1989/> (accessed on 15 December 2025).
- Plain, Anne, Rebecca Eynon, Isis Hjorth, and Michael A. Osborne. 2022. *AI and the Arts: How Machine Learning Is Changing Artistic Work*. Report from the Creative Algorithmic Intelligence Research Project. Oxford: Oxford Internet Institute, University of Oxford.
- Prada, Luis. 2025. AI Music Boss Says Musicians Don’t Actually Like Making Music. *Vice*. Available online: <https://www.vice.com/en/article/ai-music-boss-says-musicians-dont-actually-like-making-music/> (accessed on 15 December 2025).
- Roddy, Stephen, and Brian Bridges. 2026. Cybernetic Resurgences: Machine Music Be-yond AI Slop. In *Artificial Media*. Edited by Nelson Zagalo and Damian Keller. Cham: Springer, pp. 95–113. [CrossRef]
- Rovan, Joseph, and Vincent Hayward. 2000. Typology of Tactile Sounds and their Synthesis in Gesture-Driven Computer Music Performance. In *Trends in Gestural Control of Music*. Edited by Marcelo M. Wanderley and Marc Battier. Paris: IRCAM.
- Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. 1986. Learning representations by back-propagating errors. *Nature* 323: 533–36. [CrossRef]
- Schedel, Margaret, Phoenix Perry, and Rebecca Fiebrink. 2011. Wekinating 000000Swan: Using Machine Learning to Create and Control Complex Artistic Systems. Paper presented at the International Conference on New Interfaces for Musical Expression, Oslo, Norway, May 30–June 1. [CrossRef]
- Sharma, Abhishek, and Keith M. Goolsbey. 2019. Simulation-based approach to efficient commonsense reasoning in very large knowledge bases. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence (AAAI’19/IAAI’19/EAAI’19)*. Palo Alto: AAAI Press, vol. 33, pp. 1360–67. [CrossRef]
- Sixtus, Mario. 2025. AI—Death of the Internet? (TV Documentary) ZDF/ARTE. Available online: <https://www.imdb.com/title/tt38341790/> (accessed on 15 December 2025).
- Snelling, Grace. 2025. Apple TV Made Its Colorful New Logo in a de-Lightfully Old-School Way. *Fast Company*. Available online: <https://www.fastcompany.com/91437623/apple-tv-new-logo-practical-effects> (accessed on 15 December 2025).
- Stahlberg, Felix. 2020. Neural Machine Translation: A Review. *Journal of Artificial Intelligence Research* 69: 343–418. [CrossRef]

- Upton-Clark, Eve. 2025. 'It's Not Enjoyable to Make Music Now': AI Music Platform CEO Is Under Fire for Going After Human Creativity. *Fast Company*. Available online: <https://www.fastcompany.com/91264510/its-not-enjoyable-to-make-music-now-ai-music-platform-ceo-is-under-fire-for-going-after-human-creativity> (accessed on 15 December 2025).
- Vardi, Moshe. 2012. Artificial intelligence: Past and future. *Communications of the ACM* 55: 5. [CrossRef]
- Wahde, Matthias. 2016. *Introduction to Autonomous Robots*. Goteborg: Department of Applied Mechanics, Chalmers University of Technology.
- Walsh, Shelley. 2025. Timeline of ChatGPT Updates & Key Events. *Search Engine Journal*. Available online: <https://www.searchenginejournal.com/history-of-chatgpt-timeline/488370/> (accessed on 15 December 2025).
- Wichern, Gordon. 2017. What the Machine Learning in RX 6 Advanced Means for the Future of Audio Repair Technology. Available online: https://www.izotope.com/en/learn/what-the-machine-learning-in-rx-6-advanced-means-for-the-future-of-audio-repair-technology?srltid=AfmBOorLe0oQxxgoXKYXte0kot8zi-Z1SJ2YdkLaKQ_4HhuEUHsm8Ozc (accessed on 15 December 2025).
- Wikipedia Contributors. 2025. Tool. Wikipedia, The Free Encyclopedia. Available online: <https://en.wikipedia.org/w/index.php?title=Tool&oldid=1323510580> (accessed on 15 December 2025).
- Wilkins, Dennis. 2017. PG Music Band in a Box 2017. *Sound On Sound Magazine*, May. Available online: <https://www.soundonsound.com/reviews/pg-music-band-box-2017> (accessed on 15 December 2025).
- Xu, Ziwei, Sanjay Jain, and Mohan Kankanhalli. 2024. Hallucination is inevitable: An innate limitation of large language models. *arXiv* arXiv:2401.11817. [CrossRef]
- Zivanovic, Aleksandar. 2005. The Development of a Cybernetic Sculptor: Edward Ihnatowicz and the Senster. Paper presented at the ACM Creativity and Cognition Conference 2005, London, UK, April 12–15. [CrossRef]
- Zuhair, Muhammad. 2025. Micron Is Abandoning Consumer SSDs & RAM; Crucial Products Are Being Killed Off as the Company Shifts Everything Toward AI. *Wccftech.com*. Available online: <https://wccftech.com/micron-is-abandoning-consumer-ssds-ram/> (accessed on 15 December 2025).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.