



Musicking with Algorithms: Thoughts on Artificial Intelligence, Creativity, and Agency

31

Palle Dahlstedt

31.1 Introduction

In this chapter, I present a pragmatic, critical, and sometimes speculative view of what Machine Learning (ML) and Artificial Intelligence (AI) bring to the table for art and music. It is pragmatic in the sense of analyzing what can actually be done today by musicians and composers working with AI, and what is missing in terms of creative agency. How does AI relate to other technologies in the context of art? Yet critical about the popular expectations of AI, its ascribed abilities and agency, and how AI is written and talked about today in terms of creativity. No, computers cannot paint like van Gogh or compose like Bach. What is really the role of humans, as designers, programmers, users, and tweekers, behind current AI applications? Still, I try to be visionary about the long-term future of AI in art and music. Will we ever see autonomous AI artists, composers, and musicians? If so, why would they even care to make art and music for humans?

I will primarily talk about two main categories of algorithms: statistical Machine Learning (e.g., neural networks of various kinds), and Evolutionary Computation. These two categories are both wide and diverse and encompass most of today's applied AI. They share the properties that they may work with data on a higher abstraction level, find solutions to problems, and generate material of different kinds, without the specifics of these solutions or material being described in detail. They can be applied in many different ways in relation to artistic creative processes. As there is no common term for these different algorithms as a group, I will in the following use the term *AI algorithms*. When I refer to Machine Learning algorithms specifically (excluding Evolutionary Computation), I will use the term *ML algo-*

P. Dahlstedt (✉)

Department of Computer Science and Engineering, and the Academy of Music and Drama,
University of Gothenburg and Chalmers University of Technology, SE-41296 Gothenburg,
Sweden

e-mail: palle@chalmers.se

© Springer Nature Switzerland AG 2021

E. R. Miranda (eds.), *Handbook of Artificial Intelligence for Music*,

https://doi.org/10.1007/978-3-030-72116-9_31

873

rhythms. I will speak about these techniques from a more general viewpoint, and some things may not be applicable to or relevant for all kinds of algorithms. I hope the reader has an understanding of this necessary simplification.

AI algorithms can be made into powerful tools that allow for new ways of working, but they are not miracles—they have constraints and limitations, and being aware of what these are, and what the implications of using them are, is crucial for an artist. Working with a tool without awareness of aesthetic implications, or maybe without even being aware that there *are* aesthetic implications, may lead to reduced independence, unconscious shift of agency from the artist toward the toolmaker, and to artistic output that is very similar to that of other users of similar tools.

In this text, I will primarily discuss the aesthetic and philosophical implications of using contemporary high-level AI algorithms in compositional and improvisational work. I will look at how such algorithms mediate agency through the influence on the aesthetic results and also speculate on the idea of art by autonomous AI, if and when that would be possible. As my own main artistic practice and training is as a musician and composer, I will use music-making as my main example, but many of the observations are applicable also to other art forms, as the reasoning deals with the artistic creative process and creative agency in general, independent of genre. The role of technology in general, and of artificial intelligence in particular, varies only in nuances between art forms, and ideas of creative agency in music are not much different from visual art, literature, or performance.

The reasoning and observations in this chapter are a continuation of a long personal investigation of these issues, in dialog with other researchers and artists, which has involved the development and long-term use of various generative systems for music-making [18–21, 30, 31] and related more philosophical and aesthetic investigations of their implications for the creative process and aesthetic implications. Primarily, the discussion about the role of tools and agency is a continuation of my previously published theoretical work on artistic creative process [22, 24, 28, 29], and my critical view of AI creativity continues the thoughts put forward in a recent paper on big data, AI, and creativity [26]. There is only space for brief summaries of this work here, and I would refer to the original texts for a more detailed view.

31.1.1 AI and Art

There are many ideas about what creativity is, and also many different definitions and variations on AI. A great variety of algorithms have been applied under this umbrella throughout the years, but the latest AI boom has centered around multi-layered neural networks, and another important category is evolutionary algorithms.

AI algorithms can perform many different tasks or sub-tasks, such as classification of arbitrary classes of objects, outlier detection within sets, and evaluation according to trained or specified criteria, and help with decision making. They can do different kinds of optimization, e.g., with respect to similarity (hence, imitation),

computation and speed, or cost, but also with regards to more or less formal/explicit aesthetic criteria, with more or less open results. They can also optimize toward meta-aesthetic criteria, such as novelty or variation. AI algorithms can be used for predicting the likely continuation of a sequence, based on example sequences, and hence be used to directly generate output.

Such tasks, as performed by AI algorithms, can be applied to simulation, strategy, imitation, and game-play or to design, improvisation, and creativity. During the last decades, there have been numerous examples of applications of AI algorithms to tasks within the musical crafts, such as instrumentation, harmonization, and voice-leading. But this is not where the main creativity of musical composition lies. They are still very interesting challenges, similar to solving a game or puzzle (see, e.g., [56]) and an important part of composing music. The idea is supposedly that if we can solve such simpler tasks, we can go on toward the larger tasks of composing whole musical works.

Here, I will concentrate on when AI is applied to the more fundamental creative tasks:

- What happens when AI algorithms are applied to generate, suggest, evaluate, continue, expand, vary, or imitate musical material?
- What are the complications of using AI to generate music, related to the implications of training AI algorithms on existing music, and the general problem of getting AI to generate something it has not seen or heard before?

At the heart of this is a tension between optimization and exploration. Optimization can be defined as doing the best thing, the right thing, optimizing the outcome of some actions, or finding the best solution to a problem, under given constraints. Exploration is expanding the limits of what has been done before, searching the space of the possibilities for new and interesting solutions or material, or creating something different from everything seen or heard before, in a fundamental or conceptual way (not just tilted a little bit). It is not difficult to create a slight variation of something, and novelty is easy—any output based on chance operations will be novel in some trivial way. But it should be novel in an interesting and meaningful way, or at least in a way so that the receiver can ascribe meaning to it. Most AI algorithms were designed for optimization, but some are also applicable to exploration.

We will look at the implications of using AI algorithms in these contexts, the implications for aesthetic results, and for the agency.

31.1.2 Motivation

What does it solve to analyze where the agency lies? It does not make the systems smarter, and it is not (at least not primarily) about authorship or about giving credit, but about understanding. It helps us understand what the contributions are from each part of a system that we perceive as creative, and to appreciate to what degree

everything is a part of a system. It helps us talk about it in proper terms, not ascribing intelligence and creativity to a machine in those cases when the actual creativity comes from human researchers, engineers, programmers, and algorithm users, while still acknowledging the (potential) contributions of the algorithms. And it will help us realize when an algorithm really is creative.

The conclusion may very well be that it is impossible to tell where the creative agency lies. In that case, the investigation has taught us to be more humble in relation to such systems and to be careful with how we talk about them, and it may make us realize that emergence is a powerful thing.

In the current debate, it is common to see popular science articles about how AI systems are creative, how they have composed new hits in the style of The Beatles [17, 46], or how they have created pictures in the style of van Gogh [52]. In the press, it is often spoken of as if the software algorithm, the AI, has created these aesthetic artifacts all by itself. This is of course not true. Usually, the generative AI system behind such news is nothing more than a sophisticated transformation tool or mash-up engine, and there are humans behind at all stages in the process. So many design choices are taken along the way, and so much information is flowing into the AI implementation from humans, that it is simply deceptive to talk about it as “created by machines”.

There are many problems with such unrealistic descriptions. It gives artists and listeners the wrong idea about authorship and about the abilities of AI. It neglects the extent of the human agency, and while the end credits of a major movie may be exaggerated in their detail, a more adequate understanding of the attribution of agency in creative processes that involve significant generative computation may help acknowledge the influence of toolmakers and algorithm designers, mediated through algorithms, in terms of influential agency.

Sadly, such sensational attribution of agency to AI shape the general public’s expectations of AI and its current capabilities. This is not only a problem of expectations, but one of politics, ethics, and a Public Relations problem for AI—depending on your position in relation to it. While these are wonderfully capable algorithms that we should absolutely use and apply in artistic contexts, we should also have realistic expectations, and talk about them in correct terms. An important part of this is to ascribe agency to the correct parts of the system and not neglect the human agency that is (still) such an important part of all AI systems and will most likely continue to be so for a while.

31.1.3 Properties of an Artist

To get some leads on creative AI, we may ask the question: What properties are needed from an AI composer or artist for us to regard it as an artist in its own right? It is clearly not enough that the output of an AI has properties similar to the output of human composers and musicians, so we could start to answer the question by thinking about what are the properties of a human artist.

Art is not instantaneous—it comes out of processes, of varying durations. During this process, the artist interacts with her surroundings. There is continuous input, in the form of a stream of impressions and social interactions, and there is similarly continuous output in the form of sketches, temporary results, dialogs, and social interactions triggered by reactions to her art. This is a feedback loop around the creative process and the artist, with information flowing in all directions.

An artist has something to say, consciously or unconsciously. Values and views held by the artist will be there, embedded in design decisions, whether she likes it or not. If the artist does not intend to say anything or does not think that her music has a message, receivers (listeners) will read something into it. The output is appreciated (ascribed a value), and it is relatable, at least in projection. Sometimes some effort is required. The artwork conceptually relates to the world and to previous art. It is also in itself a part of a long-term process and discussion about what art is and can be.

A listener can also empathically relate to music and music-making, as many have some experience of playing an instrument, or at least of singing. We can perceive effort and intention in others' behavior such as playing music [25] also when we do not have sufficient domain knowledge to understand in detail what is going on. For example, when listening to an ensemble of improvisers, we can empathically perceive their efforts, interactions, and struggle.

So, why should we not expect these properties from an artificial artist, and from art and music created by such agents? It is not enough that they imitate output from human artists, as this is fundamentally non-creative (a process of optimization instead of exploration). We must expect new material, derived from its own actions and interactions, that is meaningful in relation to its surrounding world and its place in it. And, for it to become meaningful to us, references should exist also to our world, or such relations should be possible to form or construct for us when listening.

31.1.4 Possibilities with AI in Art and Music

We can see three main ways in which AI can be used in artistic creativity:

- As a **tool**: It can be used as a **black box** system operated by a human artist, to generate a batch of output that can be used in various ways by the composer, at various stages of the creative process, e.g., a generated sound to include in a song or score material to be further manually edited or arranged into a composition.
- As a part of a **system**: It can be part of an **interactive** system consisting of both machines and human agents, which is used to create art and music.
- As an **autonomous agent**: It can form an **autonomous** system that creates art without any interactions with human agents.

Some AI algorithms, and this is especially true for evolutionary algorithms, can help us explore what is possible under certain well-defined constraints, in that they

perform a structured search of the space of the possibilities. We may search the same space of possibilities as before, but with the help of such algorithms, we can search it more efficiently. The word efficient may not ring well in artistic contexts, but it can be understood as meaning two things. First, we can reach similar results as with traditional tools, but faster. Second, and more interesting, is that we can spend the same effort and time we would have spent with a traditional tool, but as the algorithm helps us reach new corners, farther away in the space of possibilities, we attain new results and new artistic expressions with similar efforts.

It is not only possible to reach new remote corners, but for time-based or linear arts such as music, the search path is also interesting. New tools allow us to find and follow new trajectories in the search space, and these paths shape the narrative [22].

AI algorithms allow us to work with new higher abstraction levels, in several ways. Neural networks, and especially Deep Learning algorithms, are able to learn, process, and reproduce patterns and stylistic properties of musical material, and with Evolutionary Computation, the use of high-level analysis in fitness functions can allow for control of complex properties in musical results without the need to explicitly formulate methods for generating them.

We can work with stylistic patterns in appearance (material patterns) or behavioral patterns in interaction. We can interact with algorithms by example, e.g., train an algorithm with examples that we want it to be influenced by.

Essentially, this possibility of working with non-precise and high-level input relieves the computer music composer from the need to understand and interact with code in a procedural way, from having to adjust detailed technical parameters, and from having to specify things explicitly in detail. She can instead concentrate on ideation and communicate with the algorithm through music (as done by, e.g., David Cope [15], and in any AI implementation that uses a musical training set). Still, she needs to understand and form the experience of working with these tools on this new abstraction level. It is not less complex nor less complicated. It is different.

AI algorithms can, due to their ability to accept high-level input and generate high-level output, be part of systems of connected interacting nodes, machines, or humans, and many different algorithms can work together. Each node contributes something to the overall creative process, but in such systems, it can be very difficult to say which part contributes what property of the output. It emerges as a systemic property of interacting parts, that each co-create the music [12, 13]. Thanks to the possibility of high-level input and output, AI algorithms can, just like human musicians, be designed to take their own output as input, and thus become complex feedback systems. This is a natural development, as many AI algorithms are in themselves already set up as feedback systems (Generative Adversarial Networks, Recurrent Neural Networks, evolutionary algorithms, etc.). Communicating (in a way) through actual musical material makes it easier to think of AI algorithms like human interactors, like a composition teacher, or a musical friend. Still much more stupid than a human, but in other ways smarter, faster, or more efficient (here is that word again). And, primarily, still different from a human.

This ability to work with higher, less precise, abstraction levels may allow for the composer and sound designer to think in potentials, using the definition of spaces of possibilities as a compositional design strategy. She may think in terms of what kind of textures or sonic features would be possible—and then let an algorithm explore that space. Already the definition of the space contains important aesthetic choices as input from the composer, and if the exploration is interactive, even more input is provided, e.g., by picking desired results from a large set of outputs or providing feedback underway. One can compare such a process with gardening. You choose what to sow, you tender it, tweak parameters during the development and growth (add nourishment and water, see that there is enough light, apply pruning and selection), and then you harvest the results (this analogy was further elaborated on in [26]). I often use this approach when composing or doing sound design with the virtual modular synthesizer Nord Modular G2. I know what kinds of phenomena I am interested to include in the patch, such as a certain kind of gestures, a certain kind of potential timbres, certain potentials for cross-modulation, a few filters, and some potential feedback paths. I make a patch containing all these building blocks without even once listening to the patch. When it is complete, I start exploring the parameter space of the patch using the built-in interactive evolutionary tool [18, 21]. While occasionally adding some manual parameter edits or some slight adjustments to the patch, the interactive exploration is the main driving force. I usually find and harvest lots of sonic material in this way, for later use.

A certain generation of computer music composers (like myself) learned AI by writing our algorithms from scratch. But the sophistication of today's AI algorithms has surpassed the point where this is realistic, and young people today learn to use high-level tools directly, as available in programming libraries or end-user applications. I will get back to that, but this means embracing a certain lack of control and transferring some agency to the toolmaker, just as a violinist relies on the contributions of generations of luthiers, and a composer rests on the shoulders of musical theorists through years of training, internalizing these theories.

But it has its advantages too. Such AI users have a chance of developing a new craft, given that they invest the same time and effort with these new tools, just as those from a developer background did, but on a higher abstraction level. As everybody's time is limited, they have a chance to accumulate a larger experience database than the developer-artist was ever able to, and if used critically and with reflection, they will develop a craft of applied AI in music. This requires developing an understanding, practice, and skills. Deep technical understanding may not be needed, but systemic understanding and the formation of good cognitive models of the systems formed by artistic practice, and this takes time and effort.

31.1.5 Art in AI

Technology does not only provide tools for the creation of art, but it is also a potential medium in itself, for artistic expression. Its complexity, and the societal and aesthetic implications of AI in music and art, are sufficiently interesting for it to

also be a subject for new works, where AI algorithms in themselves can be used as a medium for expression. Artistically designed algorithms, and algorithms *as* art. Not algorithms *for* art. Here I do not refer to “the Art of Programming” in Donald Knuth’s sense of a complex and refined craft [42], but to how musical AI algorithms could be tweaked, modified, and designed in personal ways, to express new thoughts in music.

AI also allows us to work with more complex material, for example, to control complex processes such as feedback networks, with complex fitness landscapes, where the interesting or even perceptual meaningful points in the solution space are far apart or hard to find. You can use an algorithm to adjust many parameters simultaneously without knowing what they mean, and hence in simple interactive processes “play” on sound engines previously hard to control, to the extent that they were previously unthinkable as musical tools [18, 19, 22]. In the same way, AI can make complex sound generation tools and generative algorithms more accessible to non-technical users, by providing more intuitive interfaces and hiding the lower computational levels.

Music is an art form that, to a large degree, builds on creating meaning through internal (within a work) and external (between works) musical references [16, 49, 50], sometimes negotiated through expectation and surprise [39]. Music created with similar techniques often exhibit similarities and form one further level of references. High-level affinities emerge from algorithmic similarities. Each new tool provides a new kind of reference, and here AI tools, operating on higher abstraction levels such as styles and patterns, can form references similar in kind to those in (postmodern) compositions that use style and imitation as the medium of expression. New kinds of references can also emerge within sets of material generated in a single process, from the same training set or from a similar genetic representation. For example, a search trajectory from a session of interactive evolution forms a narrative of related musical material [22], related to the metamorphosis or variation composition technique, as used by, e.g., Vagn Holmboe [53] and Jean Sibelius [65].

31.2 Agency

Agency as a concept goes back to Aristotle and Hume and was originally defined as the capacity of an entity to act [61], to cause something. The causal chain should not pass through the agent, but it should originate in the agent. The concept was further developed by Anscombe [1] and Davidson [32] toward an idea of *intentional agency*—when an action is initiated willfully by a conscious agent (normally a human), and as a consequence, the agent can be blamed for that action—it was intentional. In art, blame may not be the appropriate word. As art can provoke, I often use the idea of who is the agent behind artistic provocation to sort out what an intentional agent is in art. This acknowledges a sender, an author behind the work, with autonomy and intention. Can we be provoked by an artwork created by an AI? Or will we be provoked by the humans behind the implementation of AI? This

emphasizes the importance of an author in art and music. A different kind of agency is *causal agency*, which is commonly used when talking about the agency of material things. Such agency cannot be blamed on the subject.

Agency understood as the intentional agency was originally assumed to come from a human agent, but expanded notions of agency and agent appeared already with the dawning cybernetics, later implemented into AI, where an agent could be as simple as a thermostat, a reflex agent that reacts to some condition with some action [58, 70]. The idea of a software agent emerged, first formulated by Hewitt et al. [37] and further developed by Hewitt [36], as a simple self-contained interactive entity, with internal states and acting according to a script. It may be goal-oriented, even though those goals may be very simple and the behavior predictable. It can also be arbitrarily complex.

Another approach to agency is found in Bruno Latour's Actor-Network theory, where actors can be also non-human agents. He asks, in relation to any kind of agent: "Does it make a difference in the course of some other agent's action or not? Is there some trial that allows someone to detect this difference?" ([43], p. 71) If the answer is yes, which is for most kinds of tools and artifacts, it can be ascribed to some kind of agency. Latour is clear that such agency can have different magnitude or strength, and he uses a number of words to describe this kind of agency of non-human entities: They can "authorize, allow, afford, encourage, permit, suggest, influence, block, render possible, forbid" (ibid., p. 72). Latour also mentions that a non-human entity can mediate agency over time, related to how Vygotsky [69] before him talked about tools as carriers of cultural behavior, and Gregory [35] describes them as carriers of intelligent behavior. For a good overview of the idea of non-human agency in Latour's work, see Sayes [60]. Another view on the agency of artifacts has been presented by Johnson and Verdicchio [41]. Their so-called "triadic agency" presents a more applied perspective, analyzing the agency contributions from the three agents: designer, user, and artifact, and related ethical aspects (responsibility) connected to ascribed intention as divided between these three agents. Another interesting attempt at redefining creative agency as distributed among a number of contributing (and potentially interacting) agents has been presented by Oliver Bown [12].

In the following, I will primarily focus on this widened (or simplified) concept of agency as an influence—we may call it *influential agency*, as carried by artifacts and tools.

In this text, we will assume that humans have free will and are able to take decisions with responsibility for the consequences of their actions. This is the basis for the judicial systems of most societies, so it is a reasonable assumption. Still, no human is independent of external influence. Nature and nurture make us what we are, and most of our ideas and actions are related to or derived from what we have observed in others. My personal view of what music is and can be is certainly very much shaped by my musical training and the music-making I have observed from others. But it has grown into a mix unique to me, shaped by my specific biographical details. So the music I compose will be personal and unique, but it is at the same time derivative. The sources that have influenced me, which perhaps have

a causal agency in relation to my music, are so diverse, and have been mediated in so many steps, that it is hard to pinpoint specific dependencies. What we do undoubtedly depends on what others have done before us, backward in untraceable steps.

As artifacts get gradually more complex, like AI algorithms, the border between intentional agency and influential or causal agency becomes blurred. AI algorithms fall somewhere in between these two categories of conscious actors with intentions and dead material objects subject to the causal laws of physics.

Many scholars have ascribed intentional agency to AI entities [11, 40, 55, 71], at least hypothetically, while talking about future implementations of AI. Such future AI may be closer to humans in their cognitive abilities, but this is far from the situation we have today. Today's AIs are not autonomous, they do not reflect on their own intentions, and they can't be held responsible for their behavior and their choices. Still, they are very powerful tools. But as we will see, there are humans behind all design choices.

While lacking intentional agency, the AI implementations we have today are much more complex in their behavior and in their dependencies than the kind of tools and non-human agents people like Vygotsky and Latour were supposedly referring to when formulating their theories. So, it is definitely worth analyzing to what extent these algorithms have agency, and what determines the properties of this agency.

31.2.1 Influential Agency

I introduce *influential agency* as an aspect of agency related to causal agency. If we take into account that tools can be mediators (with Latour's words) or carriers of agency, from the toolmaker, realized during tool use, it is also related to authorship. If you include responsibility, causality becomes intentionality. But responsibility requires awareness and certain control over the processes behind the causality, which humans certainly are capable of having, but which non-humans may not have. The intentions of the toolmaker were presumably around what the tool should be capable of and how it should be designed to do this, but not exactly how it should be applied. So, in the mediated agency, intentionality may not carry over, as a toolmaker cannot predict what your intentions with the tools are. It makes sense to talk about the influential agency from the tool and toolmaker, but not about the intentional agency from the toolmaker.

It may seem far-fetched to bring in authorship to the equation, just by tool use, but when tools are as complex and carry aesthetic implications on different levels of detail, with stylistic implications, they significantly shape the music. A clear example is the complex tools of today's electronic music. You can often hear which algorithms, instruments, and modules have been used in a certain piece of music, and toolmakers certainly shape the aesthetic trends with their designs. In the case of AI algorithms, which can embody patterns and behavior from existing music and music-makers, this phenomenon may be even more significant.

Influence is not just brought by tools, but also by many other sources during the creative process [24]. As a thought experiment, we can recall Latour's question: What difference does a certain actor make? What would happen if we altered details in the creative process, such as swapping a tool for another, swapping an algorithm for another, or change the person who carried out some part of the process? Or change any other constraint? What and how large differences would these changes make to the resulting action or output? If this tool was different, how would the results be different? Perhaps there is also another kind of agency, related to influence, but more drastic: Without this tool, this result would not even be possible—we may call that *conditional agency*, and both can exist in parallel.

Influence brings meaning and references to the work, regardless of intentions. Tools bring references to other music composed with similar tools, that exhibit similar structure not connected to any intention from the toolmaker. Tools may give rise to internal structure in a piece. And they imprint traces of the process of coming into being of the work.

For example, a certain time-stretching algorithm (e.g., FFT-based stretching) produces a certain kind of artifacts that are easily recognizable (a kind of harmonic bubbling). As a result, when applied with extreme settings, we are listening more to the algorithm than to the original stretched sound. Each composer who applies this algorithm ends up with quite similar results, and even though these results are quite complex in structure, they are similar, and the input from the composer is small: A source sound (that is not quite audible anymore), a few parameters for the tool, and the decision to use this particular tool. In this case, the influence of the toolmaker is much larger than that of the composer.

In a similar way, a certain AI algorithm can—even though its output can vary widely—bring a certain kind of characteristic structure, or a kind of sounds, depending on what it is capable of representing or generating. The specifics of these implications are dependent both on the design of the algorithm, the particular implementation, and how its parameters are set.

31.2.2 Influential Agency of an Algorithm

The influential agency can exist in the form of mediated agency from humans.

What comes from the algorithm, and not from the human designers and operators behind it? It is not a simple question to answer. What influence cannot be referenced back to a human? Does it have to exhibit true emergence for that to happen? And how do we know when that happens?

While the human influence is certainly there, the algorithm may induce an influence of its own. There are some situations when we could expect this to happen.

First, when there is emergence happening in the system. Emergence is defined as high-level behavior that is not directly traceable to the low-level parts, e.g., the complex behavior of an ant farm, which is way more complex than each ant. Complexity makes emergence untraceable.

Second, when there is a layer of independent learning in the system, learning during the process, in interactions with its environment. Just as a person is shaped by nature and nurture, with complex untraceable influences, an AI algorithm doing this may exhibit autonomy and perhaps agency.

Perhaps it is easier to go back to our original question: What does not come from the algorithm itself, i.e., what comes from the humans involved? Then what is not included in these answers will emanate from the influence of the algorithm. Emergence is tricky, and a lot of human decisions go into it. As an experienced designer of complex systems (co-evolutionary, cell automata, feedback systems), I know that I can design systems that give rise to desired emergent results (sometimes after a few design iterations, but still), and that this is a craft that can be learned, even though it is hard to verbalize this knowledge. So human agency may go into such systems too.

So what is the human influence? In a typical AI implementation, there are many stages of human influence. The choice of the training set, parameter settings, feedback from the human evaluation that goes into design choices, data representations, tweaks during implementation, changes, decisions about the workflow of the algorithm, how it will interact with its user, and many more.

Time has a role in moderating influence. AI systems are still usually run in computation batches over a limited time span, with human interactions before, after, and sometimes during the process.

As systems get more complex, and the learning continues over a long time, say many years, with continued interactions with its surrounding, including humans, it will be harder to speak about the agency of specific humans or specific human decisions. As the number of interactions grows, influence becomes diluted and harder to trace, and as the system grows (due to learning), complexity increases which makes emergent behavior harder to explain.

31.2.3 Influence as Information

Can we define agency in terms of information, in the meaning of Shannon [62]? I will not even try to make a formalized theory for this here, but I will use this as one way to reason about agency in artifacts resulting from human and AI creative processes.

Agency is often defined in terms of action and causality, and the term originally comes from action theory [1, 32]. But an action can also be thought of in terms of information flow from cause to effect. This idea is not new and was an important idea of cybernetics already from the start [3], emphasizing the connection between information and control in both humans and machines. This was further formalized by, e.g., Touchette and Lloyd [67], and Shannon-inspired approaches have also been applied to learning [8] and decision and action [66].

In this way, the influential agency can be thought of as a transfer of information from the influential agent to the musical result. For example, adding a note to a composition adds a certain amount of information: timing, duration, pitch,

dynamics, etc. Every choice in creating music with a modular synthesizer contains an inflow of information, such as selecting which modules to include in a patch, which patch points to connect with a cable, timing information about when connections and parameter changes are done, and physical input into gestural interfaces, which are continuous signals at high resolution containing a large amount of information, which can be measured in a crude way as the size of a MIDI file containing a recording of it. Any application of a predefined generative process with a certain number of parameters introduces a certain amount of information.

It has already been mentioned that the influence can vary in magnitude [43], and perhaps we can understand this in terms of how information flows from various agents into the creative process, and how some of that information is lost in the process, and what remains. This information flow cannot always be traced backward. Some information is lost, some is redundant, and some is transformed. But there is a clear correlation, a dependency, between the information flowing in and the information contained in the finished work.

The main idea here is that a piece of music contains a certain amount of information. This information was introduced from somewhere during the creative process that led to the existence of the music. If composed by a human, it comes from years of musical training, from external impulses during the creative process, from the tools used (such as music theory or sound processing tools), from incidental actions, and from decisions taken during the process, influenced by a myriad of factors. If the piece is generated by an AI, i.e., a computer program, the information results from the processing of other information, either previously stored, as contained in the program or as input during the running of the program.

If an AI were an isolated entity with no information flowing in (except at its creation), the expressive power of the algorithm would be saturated after a certain time. If the output is only based on the internal states of the algorithm, it does not form an interactive relationship with its environment, and its output will not be contextualized or have any relationship with its environment. Or, any such contextual relation will be diluted over time, as the environment changes, but the AI will not change.

A lot of information is embedded in the training set used within many ML algorithms. Some information, but probably less, is contained in the process of selecting the training set. Other information is input during the design and implementation process in the form of numerous design decisions and parameter settings, as borrowed code from existing libraries, and from the process of coding.

Even if the amount of influence could be estimated, it cannot be isolated as having caused particular features in the results, since several agents interact and the result emerges from these interactions in a way that is not possible to attribute to each one of them, e.g., a particular artist, with her personal aesthetic preferences and characteristic behavior and creative habits, interacting with a particular tool (in a broad sense), results in a unique combination. This is of course also affected by what happens in the environment during the process. This particular interaction and its unique results could not have occurred in any other way. The same tool in the

hands of another artist would lead to different results. And the same artist using another tool would also result in different art or music.

31.2.4 Influential Agency in a Typical AI Music Implementation

AI algorithms are often talked about as being able to create artistic artifacts on their own, but in reality, there are many layers of human influential agency at play. Let us start by looking at the various points of the development and application of such algorithms where this happens.

As an example, let us consider a hypothetical generative Recurrent Neural Network (RNN) network that is used to generate music in a style based on a training set of existing music by some human composer.

- It was humans who invented the general concept of artificial neurons, inspired by biological neurons.
- In the case of this specific type of neural networks, a large number of humans have been involved in the development and improvement of the underlying algorithms, over many years.
- Humans programmed a particular implementation of this algorithm, as a library usable by others. It is still a general set of algorithms and tools, to be applied to a wide area of possible situations and tasks, but also comes with a set of constraints, following from design decisions by the programmers.
- Humans chose which particular generative algorithm to use in this project, from a large set of potential choices, and which particular implementation of this algorithm, from available software libraries, or chose to implement their project within an existing development environment that comes with a set of libraries.
- Humans also chose what hardware to run it on, which comes with a set of constraints, such as computational speed, available memory, and a processor architecture more suitable for one type of implementation than another (e.g., a certain kind of parallelism).
- Humans chose how the material was to be represented to the algorithm, which can have a large impact on what can be learned during the training phase, which features are detectable in the input, and what output can be generated.
- Humans set a myriad of parameters that control how the chosen algorithm is operating: Number of layers of neurons, number of nodes, size of the training set, training parameters and sub-algorithms, preprocessing of the training set, etc. All these parameters have implications for how the algorithm will perform and usually take some experience to get right. There are usually no default choices for such parameters that work for all kinds of projects, and often considerable experimentation is needed before good settings are found.
- Humans choose the training set which has crucial aesthetic implications for the generated output.

- Humans tweak parameters of the algorithm, based on iterated outcomes, if the algorithm does not work as expected, the result is not good enough, or shows unwanted features. In this feedback process, human aesthetic evaluation is a crucial component.
- Finally, humans select the best examples from a large set of generated outputs. Also here, human aesthetic evaluation is at play.

In most cases, all these humans are different people, working at different points in time, and their decisions were taken in very different contexts, at varying distances (in time and computational steps) from the final artistic result. Some of them were working in a very general and abstract context, not even considering what these algorithms will be applied to. Still, their decisions carry an impact on the final results. Sometimes the humans were part of the feedback process of trial and error or final tweaking, trying to make the final results as good, or as similar to the intended style, as possible. The decisions, choices, and values of those early in the chain remain embedded in the tools they develop, which are then used by others in the later steps of the chain. All these steps have an influence on the output.

31.2.5 Influential Agency in an Actual Example: Ossia

Let us look at a similar analysis of a real case. I have chosen my own Ossia system which is an implementation of an evolutionary algorithm that breeds complete performed piano pieces of a duration of 30–90 s each. It started as an interactive evolutionary composition tool in 2000 [19], and autonomous evolution (based on random or keyboard input) was added in 2002 [23] when it was exhibited as an installation for computer and Disklavier player piano at the Gaudeamus Festival in Amsterdam. Later it was exhibited for several years at the Universeum Science Center in Gothenburg, Sweden. It was chosen as an example because it is well documented, and I know it inside out, as I designed and coded it from scratch in C++ in a number of versions over several years. It is still shown sporadically in lectures, e.g., as background music to the lecture versions of this chapter. Ossia composes a continuous stream of new piano pieces/performances and performs them as a suite. It can be interactive, but I will here talk about the autonomous version, as it appears to be composing by itself, and the output is quite varied.

- The choice of algorithm was influenced by my previous work on interactive evolution for sound synthesis [18], which was in turn influenced by my reading of introductory literature—introduced by my doctoral supervisor—on the topic of Artificial Life, and by my previous extensive experimentation with random search in sound synthesis parameter spaces. The idea of simulated evolution in computers goes far back and was initially mentioned by Alan Turing and John von Neumann, further developed by a number of researchers in the 1950s and 60s, and popularized as the genetic algorithm by John Holland [38]. My understanding of evolution was also influenced by reading Darwin and Dawkins.

All these sources indirectly or directly influenced my implementation, mediated in several stages through researchers, authors, and teachers.

- The genetic representation of Ossia was designed by me, influenced by knowledge of and previous work with tree-based data structures (from computer science studies and other machine learning experiments) and generative grammar (from various books and from knowledge of Lindenmeyer systems), and my previous works with recursive algorithms (e.g., fractal graphics) and recursive programming (e.g., Prolog). The idea of recursive pointers was also based on the idea that a core property of the musical form is repetition with variation, and this construction made this possible. The modifiers (of velocity, duration and pitch) were one way to make possible another desired set of archetypal musical features: exponential crescendo/diminuendo, exponential accelerando/ritardando, and repetition with transposition. These ideas were most likely influenced by a lot of reading of classical form and music theory and of musical aesthetics.
- The genetic operators were chosen and designed by me, influenced by various papers on mutations and cross-over in tree-shaped genomes.
- The initial population can be either randomly generated trees, a set of arbitrary musical sequences (stored as MIDI files, parsed into genome trees by the system), or a set of previously evolved musical pieces (in stored genome form). In the first case, “randomly generated trees” are not entirely random. They are generated in a detailed process coded by me, with choices for what random distributions there are for certain properties to appear in the tree, and what the ranges of these values could be. So it is still much influenced by aesthetic choices by me. In the second case, the set of musical “seeds” has varied, but the most used set has been a collection of simple archetypal musical gestures such as an upwards scale, an arpeggio, or a repeated note. These simple seeds were used because I thought they would expose the kinds of variations that the system was capable of. Even though this was the initial educational motivation, the set of seeds have been kept for most performances with the system. Their exact form and selection were certainly shaped by my long training as a classical musician. In the final case, there is a procedure coded by me for selecting when a previous musical result will end up in the seed pool.
- The workflow of the system, which is behind how it acts and interacts with the world around it, was designed by me based on a series of circumstances. The initial interactive evolutionary algorithm was designed for it to be used as a composition tool, generating raw score material to be arranged into compositions (e.g., my chamber work KARG [19]). Soon after, in 2002, I made it evolve pieces of its own, influenced by an opportunity for a performance at the Gaudeamus Festival in Amsterdam and by the availability of a Yamaha Disklavier concert grand player piano. It was exhibited with this piano, and I thought it would be interesting to also let visitors perform on the piano, and let the system evolve further based on the human performances. If nobody plays, it evolves new pieces from scratch.
- The fitness criteria were designed by me, in a design process that extended over several design iterations. In the first version, the system was only interactive, and

all examples were auditioned and selected by me. When I made it autonomous, I used my observations and notes from how I tended to select musical examples and tried to implement hand-coded formalized versions of the same selection criteria. They were based on statistical measures such as note density, tessitura, information content and repetition, and variations over time of these statistical measures, as a way to indirectly enforce dynamics and variation (or process) within a piece. Clearly, my aesthetic preferences influenced the workings of Ossia here.

- Evolutionary parameters (population size, mutation rates, halting criteria, etc.) were first set ad hoc by me based on previous experience of evolutionary algorithms. They were then gradually adjusted based on results of repeated test runs, based on personal preferences and informal performance evaluations.
- A final selection of pieces to be played has sometimes been made by me, as in the Ossia Suite, featuring 27 piano pieces [27]. They were selected by me from a large body of output, based on aesthetic criteria. In the real-time installation version, no such selection takes place.
- During development, the system was tested using a simple piano sound. There were two primary reasons: First, I am a pianist myself and feel very much at home with this timbre. Second, the piano is regarded as some kind of “universal” instrument, mostly because it can be played over the full pitch range by a single musician. This undoubtedly biased the design toward producing musical material that works well on the piano. In spite of the supposed universality of the piano, all instruments are different, and the way piano responds to dynamic playing, how dissonances behave perceptual in different registers, its rhythmic pregnancy, and many other features are unique to the piano. Music composed for it may not work well with other instruments. This simple design choice strongly influenced the aesthetic output of the system. And indeed, the output of the system sounds similar in aesthetics to my own piano improvisations, even though there is no training involved, nor any musical knowledge explicitly coded into the system.

The above is of course only a simplified analysis of the various sources of influence on the Ossia system. There are many more design decisions involved, and many features of the system that I cannot include here for reasons of space. Still, it is clear from the above that even though it may appear to an observer of the player piano continuously performing an endless series of new compositions that the software composes these pieces, clearly there is the extensive human agency involved. As the author of the system and a listener to maybe thousands of pieces composed by it, I can certainly hear the patterns, even though it is (designed to be) very varied. Although there is infinite variation within the result space, it is not infinite in its extension. And I am starting to grow tired of it. Exploration has mapped out the limits of the result space, if not all subspaces contained therein. I can clearly see the bars of the fence it is caged in.

31.2.6 Agency is Where in the Code?

There is usually nothing specially “AI” about the programming languages that are used for implementing AI algorithms, as they can be implemented in any common programming language. Some parts of the program implement an AI algorithm, and other parts take care of the more mundane tasks—general infrastructure of the program, the main loop that controls actions, asks for new input, asks for new output, performs memory management, etc. As all AI implementations consist of such mundane tasks and rather basic mathematics, just iterated many times, and/or in large parallel configurations, it is hard to say where in such a program agency would appear. It is often considered to emerge from the sheer scale and complexity of the algorithm.

It can also be argued that many current ML algorithms, such as deep neural networks, harbor no agency if understood as a capacity to act. As the actual AI algorithm is not associated with actions, but with evaluations and classifications, the only action-related agency in such systems happens in the “normal” code around the ML algorithm: the *main* loop, the *if..then* statements acting upon the evaluations of the AI algorithm. And these parts typically do *not* learn. And as long as these parts of the code still consist of common *for..next* loops and *if..then* statements, they will never be able to act in any intelligent way. According to this view, today’s AI systems act stupidly and repetitively, but have an ability to develop and learn complex evaluations and classifications. (This particular argument was developed in dialog with Karoliina Salminen, principal AI engineer at Huawei, Finland.)

The above argument fits quite well with the Ossia system. It is programmed to compose a new piano piece as soon as the previous one is being performed, with a fixed number of seconds between each performance. If a human plays on the piano while Ossia does not play, it will evolve a new piece based on the human input as soon as the human has been quiet for a few seconds, and then perform this as soon as it is ready. If somebody presses the Q key on the computer, the system stops. The fitness criteria are varied according to a number of preprogrammed parameter sets, giving quite different outputs. There is not really any potential for long-term progression.

There is not much intentional agency in this scheme, which is basically a looping script with some external sensors (MIDI input from the piano and the computer keyboard) and some actuators (the player piano). Just like the simple thermostat, it is a reflex agent, although with slightly more complex internal states.

31.3 Tools and Humans

Most art and music are made with tools. As we have discussed earlier, tools are carriers of embedded agency and carriers of intelligent behavior. They influence the artistic results, because they define what is possible, and they define the paths in the space of the possibilities, along which creative process can travel [24]. Through this

influence, they carry their agency, in several ways. They lead to characteristic results, because only certain specific things are possible to do with a specific tool. Certain things are easy to do, and certain things are hard.

Previously, instrument makers provided simple tools to be used by skilled artists (e.g., a violin maker). Not simple in terms of construction or craft, but in terms of time-complexity in their interaction. They provide constraints, but interaction is primarily based on real-time responses to direct gestural input. No pattern comes out that is not detectable in the input.

Today, instrument makers provide tools that contain extensive databases of presets and algorithms for putting together this material with potential for creative agency. Such tools have stylistic implications, as they are designed to be used within a specific style of music. They can be used at different levels of control:

Pressing play

where the tool generates a whole song, or at least significant portions of it, for example, by combining complete loops and ready-made drum patterns, applying automatic accompaniment engines, and mixing algorithms. The song you create in this way risks being very similar to the song I will create if we use the same tool and there is little inflow of information.

Collage

where you manually put together finished pieces from a database, such as loops, drum patterns, and sequence phrases, and select the sounds they play. Here, more effort is put in by the user.

Detailed

where you have control over every parameter of sequences, sounds, and processing.

There is a qualitative difference between the old instruments and such new tools, which has been analyzed in detail; e.g., by Nilsson [54]. This difference is further amplified when you take the next step toward tools based on complex systems and AI algorithms, which have complex internal states which develop over time. When designing such emergent systems, the output of which we cannot predict, design choices have to deal with low-level behavior while the consequences appear in high-level behavior, and new skills are needed to understand and work with such systems.

31.3.1 Effort Versus Tool Complexity

We can do a simplified analysis of the role of effort and inflow of information in relation to tool complexity. To simplify, we only talk about small or large effort (time and amount of interaction invested), and simple and complex tools (designed by myself or somebody else). In this analysis, we must remember that influential agency can be mediated in two ways: by a user, through previous learning from others, and by a tool, through design from the toolmaker.

There are a few obvious cases:

- **A simple tool with little effort:** This does not lead far, and results will stay at the level of playing *Twinkle, twinkle, little star* with your index finger on an unfamiliar instrument.
- **A simple tool with a large effort,** applying skills from years of training: The main part of the influential agency will be from me as an artist, except that a significant part of my skills indirectly come from others, as mediated through learning. Very little will come from the tool, except in the form of generative tool constraints.
- **A complex tool with little effort and interaction,** for example, a tool containing presets and generative algorithms: The main part of the influential agency will be from the toolmaker because I will rely on ready-made material or material generated from algorithms with default parameters.

The following cases may be less obvious.

- **A complex tool of my own design, used with little effort:** If I have written every line of code, but I let the tool do the choices, the main part of the influential agency will still be from the toolmaker—but that is me. Still, a large part of the influential agency comes from the inventor of the class of algorithms I used, and the teachers or authors that taught me those algorithms. But there has been a significant inflow of information from me into the creative process, because I designed the tool. But there is a catch. If I continue to use the same tool to produce a large amount of musical output, without much new interaction or effort, it will converge toward the situation in the above case 2. My initial effort will fade in proportion with time, as the extensive but limited one-time effort will have been used to create an ever larger amount of music, thinning out the content in relation to the total inflow of information. In a sense, it can be regarded as me re-instantiating the same piece over and over again. This is the case with the autonomous version of my generative composition *Ossia*. It keeps generating music but does not add anything new.
- **A complex tool with large effort:** If I use a complex tool designed by somebody else, containing databases or generative algorithms of which the inner workings are not known to me, and I put in a large effort, with a large inflow of information, then I have time to form a cognitive model of the tool based on experience. This helps me navigate the pathways in the result space. Through the effort spent, I have a chance to find distant corners in the result space that may not be found by others that put in less effort. And I have a chance to find particular pathways in the result space that are personal to me.

31.3.2 Non-mediated Agency in Algorithms

If we, by definition, assign agency to the toolmaker, we risk ending in a paradox. As all AI systems are man-made (it lies in the word *artificial*), there is a possibility that at some point the system starts to exhibit agency of its own. We do ascribe creative agency to ourselves. And we have to be able to tell what is the qualitative difference, when they attain this agency and break free from us, or we end up in the “but who created us” circle of reasoning, and will keep looking for a first mover.

So, when does a system attain agency, aside from the mediated influential agency of the toolmaker or designer? Bown and McCormack [13] have defined what they call *creative agency* as the creative contribution attributable to the actual system, and added that “novelty and value that cannot be directly attributed to the computational system should have no weight in supporting claims about the creativity of that system”. When analyzing creative generative systems, creative agency is *the* important property of a system, not the actual creative output.

But what are the criteria from creative agency? When do we actually take our virtual hands off a system, and it starts creating beyond the influential agency of us as designers? It is not easy to answer this, but the following questions may help us on the way:

- What remains of our design over time—what is fixed and what is dynamic in the system, at different abstraction levels?
- Does the system search the same solution space each time or does it develop over time?
- Is the solution space searched in the same way and based on the same criteria each time or is there potential for learning?
- Are the aesthetic constraints of the underlying representation sufficiently relaxed or even open-ended?
- Is there a sufficient inflow of information from interaction with other agents or other parts of the environment?

Without having any definite answers, it seems to me that it comes down to process, and the internal changes of the system as it learns, transforms, or evolves. It needs to be able to accumulate impressions over time, and we need to allow for time for things to develop, while we remove ourselves. Given sufficient open-endedness, as a system gets more complex, and the learning continues over a long time with continued interactions with its surroundings, including humans, it will be harder to speak about the agency of specific humans or specific human design decisions. There will be many more interactions, and with a larger number of humans, so agency will be more distributed and harder to trace to specific events or agents. And the system will potentially grow in complexity, and perhaps approach the kind of emergence where underlying causes cannot be identified at all.

31.4 Spectra of Agency

A spectrum is a range of a certain quantity, such as the spectrum of audible frequencies. There are certain properties involved in this discussion about the agency that have a range from simple to complex. We will look at these spectra in this section, in an attempt to understand these parameters better.

31.4.1 Spectrum of Tool Complexity

We can think of tools as having two primary levels of complexity, even though they are related. One is the *level of abstraction of the material* that a certain tool or algorithm operates upon. Is it basic musical atoms such as individual sound samples or individual notes? Or more complex constellations of such atoms, such as phrases, patterns, or even operations on the stylistic level? The other is *the complexity of the operation*—what kind of transformation does the tool bring? Is it a simple linear transformation or a more complex operation? A few examples may help to illustrate this. A simple operation may be a transposition of a score, by moving all notes up by a major third, or changing the volume of a sound file by multiplying all sound samples by a constant. It is a simple, straight-forward operation applied to the basic material. An example of an intermediate operation could be to search and replace a given interval sequence with another sequence or to generate material with the same statistical interval and rhythm distributions as a given piece of music. A complex operation could be to compose a fugue in the style of Bach or to initiate a new musical style by creating a piece of music that differs in structural properties from all previously existing music. These operations get more complex both in terms of the amount of information that is introduced with the application of the tool—a single parameter in case of transposition and a complete set of Bach's fugues in the case of fugue generation. They are also more complex in terms of the amount of computation required.

A proposed simplified spectrum of tools could look like this, where the level of abstraction and complexity of operation have been combined into a single, rising scale (see also Fig. 31.1):

Simple tools

Straight-forward linear sequentially operating physical, virtual, or theoretical tools. Examples: A pencil, a keyboard, a violin, a pair of scissors, or cut-and-paste operations.

Template-based tools

Tools that contain predesigned databases of material or parameter sets, to be able to quickly solve complex tasks in predefined ways. Examples: Preset-based synthesizers and effects, a clip art database, or a loop library.

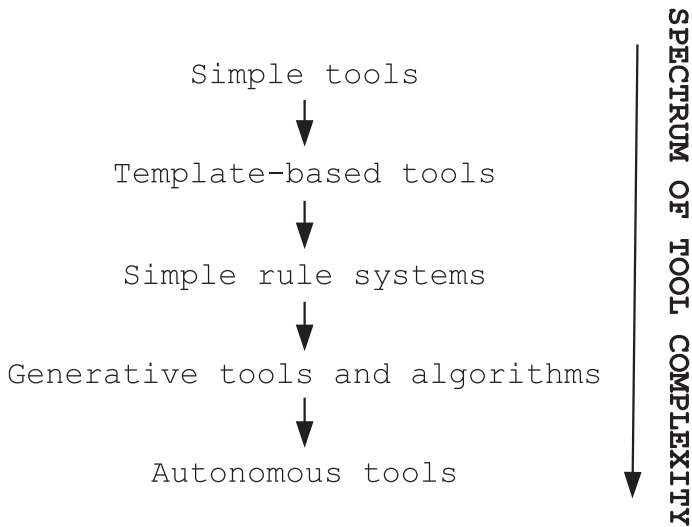


Fig. 31.1 A proposed spectrum of tool complexity. The complexity of operations increases from top to bottom

Simple rule systems

A set of behavioral rules or basic procedural code that constrain the output, and help project beyond the artist's imagination. It is easier to invent a few rules than to predict what the result will be, as we are lousy predictors. Examples: A line fractal implemented as a recursive Logo script, search-and-replace, regular expressions, isorhythmic composition techniques, a generative modular synthesizer patch, or the rules of a game piece in the style of John Zorn.

Generative tools and algorithms

More complex computational tools that generate or process material based on advanced algorithms. Possibly designed by, tweaked by, or interacted with by the artist. Examples: Most current AI algorithms.

Autonomous tools

Tools that generate or process material without any interaction with a user. If a tool becomes autonomous, it may perhaps not be called a tool anymore, as being used by someone could be considered part of the definition of a tool. No examples of this category exist yet.

31.4.2 Spectrum of Agency

Two properties shape the amount of influential agency an agent may have. The first is *the amount of interaction* with the agent or *the amount of information* embedded

in the tool that mediates the influence. The second is *the distance of this interaction or the application of this information*—in time, processing steps, computational layers—from the actual result. Information may be diluted, transformed, or lost in the process. An agent may be very distant from the final result, but still have had influence through actions, e.g., the inventor of a class of tools a long time ago, or of a new kind of algorithm, e.g., a new kind of neural network, without any intention for it to be used to create art. Such an inventor still has an influence on me making music with these algorithms, since it opened up for the possibility, created a potential. It is not intentional, but somewhat casual, and definitely influential.

Tool designer

The inventor of a tool or a class of tools.

Toolmaker

The maker of a particular implementation of the tool.

The tool itself

The tool carries embedded agency from the tool designer and the toolmaker, as a mediator. In the encounter between tool user and tool, the tool's potential is actuated in a process constrained by the tool's space of possibilities and the user's aesthetic preferences and skill [24].

Tool user/Artist

The agent who applies the tool in a specific context, to a specific material, with specific parameters. The artist also brings in influence from her own previous music, from other music and musicians, and in an extended sense from all of music history.

The artwork itself

The artwork carries accumulated embedded agency from all previous agents.

The receiver/listener

The listener, who finally receives the music (or a co-player who hears a colleague play on stage), carries out, consciously or unconsciously, an interpretation of what is heard, and through this applies influential agency. At this final stage, there is also the sense of personal agency an engaged listener can sense, as if it was created by yourself, often enhanced by dancing or other synchronous movements [25].

As shown in Fig. 31.2, there is also some influential agency going in the opposite direction. Tool users often give important feedback or requests to toolmakers, and this can be iterated many times. For example, as a musician, I have worked closely with electronic instrument makers, both in testing, suggesting new functions, and sometimes even initiating new tools. These tools contain influential agency from me and the toolmaker and are then applied again by myself. The receiver—and the artist herself, as mediated by the work—contributes to the aesthetic context, and in the long run to art history, which in turn influences the artist.

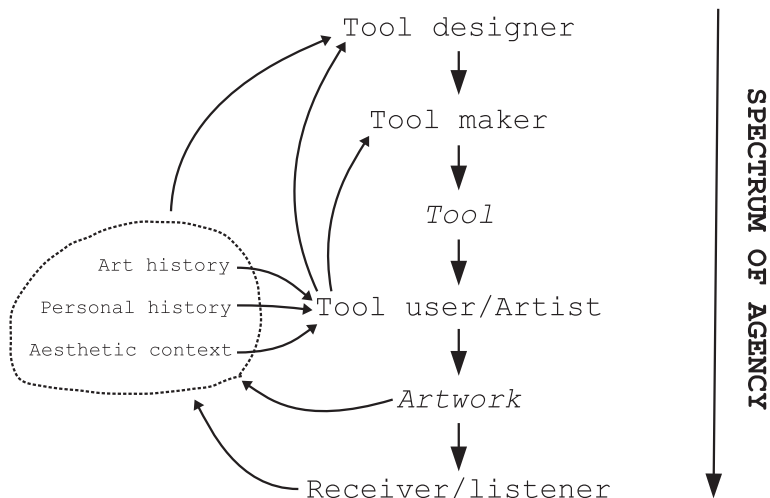


Fig. 31.2 Spectrum of agency. Influential agency flows from the tool designer and toolmaker, as mediated by the tool itself, to the tool user (the artist), and onwards, through the act of interpretation (which sometimes involves interpreters/musicians). There are also circular flows, as the artistic result becomes part of the aesthetic context these agents live in, and in the long term, becomes a part of art or music history, and hence influences both tool designers, toolmakers, and future artists

In electronic music, the tool designer, toolmaker, and tool user are often the same person. Still, electronic musicians build or program their musical tools (instruments) from other tools at lower abstraction levels. For example, a live coder may create her own language or library of functions to use in a live setting. These are usually implemented in another, more general programming language. A synthesizer builder constructs her machines from circuits designed based on decades of development by the analog synthesizer community, perhaps with added inventions of her own, and/or in new configurations, just as most pieces of music are created based upon patterns and forms that have been developed over centuries. In the same way as all music is derivative, tools are derivative.

31.4.3 Spectrum of Generativity

Based on the above spectra, we can also make an attempt at a rough spectrum of levels of generativity in music, i.e., a list of categories of music in rising order of generative complexity, and rising order of influence from underlying tools and algorithms (see also Fig. 31.3):

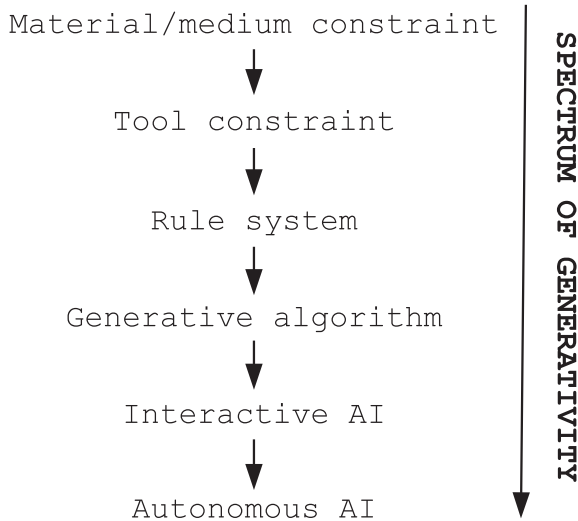


Fig. 31.3 Spectrum of generativity, with generative influence going from low level at the top to very high level at the bottom. The initial choice of medium to work in affects what can be represented. The chosen tools constrain what can be done, and successively more complex algorithms bring higher level generative properties to the creative process, gradually shifting the influential agency toward the tool designer and toolmaker—which may very well be the same as the artist

Material/medium constraint

The chosen medium has implications for what can be represented. For example, if you work with Western notation, only certain kinds of rhythms and durations are possible to write.

Tool constraint

The choice of tools has implications for what can be done in the chosen medium. Tools define pathways (structured subspaces) in the space of possibilities defined by the medium. If you work in the medium of electronic sounds, constrained by a sequenced synthesizer as a tool, certain sounds are possible as defined by the instrument/synthesizer, and certain events and kinds of parameter changes are possible to represent and control from the sequencer. Here, both the medium and the tools used have generative (and restrictive) properties.

Rule systems

Simple rules about how constituent parts interact during a creative process lead to characteristic patterns.

Generative algorithms

Computation-based algorithms generate musical material, for example, from an AI algorithm.

Interactive AI

The AI is a node in an interactive creative network of agents. It generates output based on input and based on stored/learned information. Learning may happen during the process.

Autonomous AI

The AI carries out a situated process in relation to the creation of a single work and in relation to a series of works. It has sustained artistic output, and it has a representation of and a relation to the outside world. In this situation, the AI may achieve intentional agency.

As the level of generativity increases toward the end of the list, the influential agency is shifted toward the tool designer and toolmaker, as the tools by mediation bring complex behavior and aesthetic implications that are not under the direct control of the artist. However, especially in electronic music, it is common that a composer designs her own tools.

Note also the similarity to the spectrum of tool complexities shown earlier.

31.5 Problems with Creative AI

In this section, I will briefly discuss a few fundamental problems with current AI algorithms in relation to creativity. These problems have been described and analyzed in more detail in a previous publication [26].

31.5.1 The Inherent Non-creativity of Statistical Machine Learning

Many current AI algorithms are “mean machines” that are designed to find or produce the most likely outcome. This includes most statistical methods (e.g., Markov-based models and neural networks). They will stay inside the box by definition.

Such a system learns from the training set and is then used to generate something that has very similar statistical properties. It will not be able to generate anything that adds something new, at least according to an understanding of the model as the best possible (in the given algorithm, size, etc.) representation of the complete material in the training set. Briefly, the argument is that if it is trained on, say, the set of 15 two-part inventions by Bach, it will be able to generate new inventions in the style of Bach. But each of those newly generated inventions will be completely based on what the system learned from the training set of the original inventions. And each original invention, while sharing some properties with the others, added something that was unique. If Bach would have written the 16th invention, it would have added some new ideas that were not present in the previous 15. But the AI

algorithm is not able to do that, because the impulses that led Bach to the material in his hypothetical 16th invention would have come from outside of the existing 15 inventions. It could have come from an impulse, a paraphrase on an earlier piece in the same key, or a musical idea from his wife.

If such a system would produce something that went outside of the training set, it would be thanks to its characteristic inabilities, basically its flaws and limitations that make it unable to reproduce something perfectly. This may very well result in interesting output, but it will not be because of inherent creativity, but from design faults or conscious limitations.

This relates to the previously mentioned complementary relation between optimization and exploration. Most ML algorithms are convergent and optimize in relation to criteria such as similarity. An interesting parallel—which highlights why this is so important in an artistic context—is that it is very similar to the relationship between entertainment, which aims for the middle of the circle, with known responses, and art, which tries to extend the circle, testing new ideas.

Creativity is about creating new patterns, and new kinds of patterns, while most machine techniques for handling large data are about detecting, classifying, or reproducing patterns of a known form.

Machine Learning algorithms produce models that are based on correlations in superficial observations, not on causality between interactions with the environment and artistic output. This is a classic problem in empirical science, and it is especially troublesome in this field because it tells us what but does not help us understand how and why. Such algorithms are capable of generating mimetic output, which lacks all connections to the situated process from which the originals emerged. We generate diluted Bach music, but not a model of how Bach would interact given certain input, or how Bach would have developed beyond the last item of the training set.

Similar problems appear with evolutionary algorithms when fitness functions are formulated in terms of properties of the output. This is also an optimization process where the goal is well specified. Still, there are a number of successful attempts at exploratory approaches with interactive evolution [9, 18, 69] and novelty search [44].

31.5.2 Opaqueness of AI-Generated Material

Both neural-based and evolutionary algorithms suffer from the problem of opaqueness of the results. The meaning of individual weights in Deep Learning networks is very hard to detect, and evolution often generates complex but undecipherable solutions, as pointed out already by Sims [63]. The solutions work, but we do not know how. This problem also exists in natural evolution and artificial breeding. Explainable AI has been part of significant research efforts lately [59] and is often mentioned as a necessary part of ethically sustainable AI.

This opaqueness may also lead to material that is hard to work with because we do not understand its inner logic. Even if the composer has written the code herself, the output may feel alien to her, and it will take a considerable effort to learn to work with it. For an example of such a problematic result, and how to deal with it, see [28, 29].

The flip side of this is that while solutions may be hard to understand, they may work well. When evolving or training a co-improvising agent, this may be what we need.

31.5.3 The Lack of a Model of the Outside World

The lack of semantic understanding and focus on patterns is a general problem with AI. They are applied to the recognition or generation of empty patterns lacking connection to the outside world. These generated patterns may inherit an implicit referential connection from training sets, but it is never re-confirmed and calibrated as humans do, constantly. We go back and check if an answer is reasonable if it is compatible with our cognitive model of the world.

A common rule of thumb is that model complexity should match domain/data complexity, for optimal learning without over- or under-fitting. Still, simpler models are usually preferred in AI, because they are faster to compute, easier to understand, tweak, and design, and will not over-fit. Simpler models make it impossible to represent the real world in its complexity.

The lack of world models in AI implementations enhances the previously mentioned problems with mimetic pattern generation, free from causal connections, lacking interaction, situatedness, and semantic dimensions.

Even if hypothetical more advanced AI algorithms would be able to partly represent the outside world, there is a logical paradox, developed around the concept of embedded agency [34]. The learning agent is part of the world, and hence, smaller than the world. So, a complete model of the world (including itself) is impossible, because the world is larger than itself (itself + the rest of the world). Artificial agents who realize this are subject to the same limitations and feedback loops as we are. Without this realization, it will be decisive but wrong.

To summarize, most current AI creativity imitates the consequences of creativity instead of implementing the creative process and emulating the causal chains behind it. It lacks continuous interaction with its environment and the related inflow of information. As long as this approach prevails, AI creativity will remain a mimetic black box batch process.

31.6 Aesthetics

In this section, a few issues related to the aesthetics of AI-generated music will be discussed.

31.6.1 Autonomous Aesthetics and Agency

The idea that an artwork should be judged as independent from the artist was put forward by Roland Barthes in his famous essay the *Death of the Author* [4]. After being published, the text lives its own life, loses the connection to the author, and should be evaluated as a separate entity, or as Barthes says: “it is language which speaks, not the author” (ibid, p.143).

But the artwork embeds values and patterns from the artist. It is shaped by the process of the artist and her interactions with her surroundings during that process. It is evaluated or appreciated in a context of an artist, style, and culture, by a receiver (critic, listener) who is part of that context or part of another related context. It is created with tools that are part of that context, and the tools and technology behind it are often audible or detectable for the educated listener or the fellow practitioner.

The conclusion is that the artwork does indeed relate to the world. There is an information flow in and out between the world and the creative process, and this shapes the result. So, in my view, it is a utopian view that the artwork could be regarded as completely separated from the artist. The artist, the artwork, and the receiver are all part of the world. An AI artist needs to be part of this world, too, and relate to it. It needs to be an agent in this world to create art that is meaningful to others in this world.

Part of Barthes’s argument is that the content of the text comes from numerous sources and numerous authors, and how all work thus is derivative. According to him, a text is

a multi -dimensional space in which a variety of writings, none of them original, blend and clash. The text is a tissue of quotations drawn from the innumerable centres of culture. [4, p. 146].

This is of course also true for music. In resonance with my argument above about influential agency, and mediated agency from tool designers, toolmakers, teachers, and other sources of influence, we could rephrase Barthes’s claim as: The death of *one* author, and the acknowledgment of *many*. It is in line with the idea of distributed influential agency, and a similar argument has been presented by Bown [12], in the context of computational creativity.

In a sense, the idea of the autonomous work from an analysis point of view does not take away from the fact that there has to be a creative process involving human or artificial agents, and that the conditions and properties of this process influence the result. A synthesist of art, i.e., a developer of creative algorithms, will still need to understand this mechanism of influence even if the analyst does not care.

Also, empathy and empathetic experience play a major role in the appreciation of music, in particular of live music performance, and here the human agency and our ability to perceive intention and agency play an important role. One could go as far as saying that in some musical contexts, for example, in free improvisation, the sonic result is not the most important part of the experience—that may instead be the empathetical experience of the interactions, the efforts, and the struggle of the musicians [25].

31.6.2 Characteristic Inability

In my previous writings on creativity (e.g., [24]), I have talked about characteristic inability or characteristic incompetence as the personal way in which an artist or musician cannot do what is considered “perfect” or what originally intended. I cannot write a perfect Bach fugue, and the personal way in which I fail becomes my personal style of fugues. These peculiarities are what makes it a Dahlstedt fugue. The same applies to generative algorithms—the way a generative system does not generate something perfect becomes its characteristic “personal” style.

Something related to pareidolia comes into play here. We read consistently “clumsy” results as personal, while anything perfect is not interesting (except for that virtuosic “awe”). But we are attracted to in what way it is not perfect, to the relation between perceived intention and actual result. Our own aesthetic sense sees the intended (ideal?) image behind the limited depiction and is able to abstract the transformational layer of imperfection, and this becomes an aesthetic experience of its own. The imperfections that emerge during the implementation process also adds ambiguity, hence it becomes open for interpretation and projected complexity. My reading is different from your reading.

31.6.3 Apparent Agency Attribution

Agency attribution is affected by many factors, e.g., synchrony, anthropomorphism, convention, and expectations. We see meaning and patterns where there are none—our wonderful disposition towards *pareidolia*—which makes abstract art a dangerously comfortable field for generative art. In the same way, we see agency when there is little of it, as we do not have the whole picture and we want to see an agent behind something—it makes what we see conceptually coherent. Hidden actors are harder to imagine and are easily forgotten or ignored.

Computer music and computer art are often presented to be reflected upon *as if* they are made by a human, and the degree to which they succeed is commonly used as an evaluation criterion in AI art, often presented as a version of the famous Turing test [68]. Typically, a number of musical examples are played to a set of listeners, who are asked which they think are composed by a musician or by a machine. As a number of practitioners and researchers have pointed out [2, 6, 7, 10, 57], this is a questionable method, or at least a questionable name for it, as the original Turing test, as described by Turing himself was a test of the quality of interaction with another agent, and a way to try to judge if it was conscious or not. Now, the output is instead evaluated in terms of human-likeness. As a program can contain large amounts of stored information (as embedded learned patterns or explicit material), judging finished music says nothing. In generative algorithms, evaluating variation over a large set of output or, where appropriate, interacting with an AI musician may approach a Turing test for music.

Does the agent of art matter? Humans can show a strong admiration for abstract complexity, as found, e.g., in mathematics and physics visualizations. Many also

have admiration for natural beauty, with no intentional agent behind it. Some project an agent (a god or something else), while others, like myself, see emergence from a complex system—most scientists agree that evolution and physics are not guided by intention.

Also, as mentioned above, some regard the work and author separated (autonomous aesthetics). Then, having no clear agent would not make a difference for the receiver. But can the work as a complex entity be reflected upon in different terms than how it was conceived, who did it, why she did it, what are they trying to say, and what happened in the creative process? An artwork tells me about its coming into being, and that is the meaning of the process. And my reflections and personal associations from it, even if disregarding or honestly ignorant about the author, will be influenced by this process, as the work was shaped by it.

We could also consider AI-generated art that presents itself *as if* it had a process, even if it did not have one, or if the apparent process did not correspond to its actual genesis. This would require a representation that takes such deeper layers into account to be able to coherently fake traces of a creative process.

Consider a different hypothetical case: We know the music was composed by an artificial agent. Does this devalue the art or increase its value to us? I am not talking about monetary value here. There may currently be a certain wow factor and some awe, as we are impressed by the novelty of what is possible with machines. There is also the computational sublime [48], the aesthetic awe of what is possible with computing. This may disappear if we believe it is made by a sentient being or a human (even if it is not) and is instead replaced by an empathetic perception of the efforts and skill behind the work.

This awe is relative. That which goes beyond our understanding is obviously relative to what we can understand and relative to our cognitive models of what we are trying to do. The play with this border is what we do when creating generative art. Our predictive capacity and its limits are crucial. As new generations grow up with AI-generated art and music, this may change considerably, just as how our current perception of generative computer music would be inconceivable just a few decades ago.

What about the effort in machine-composed music? Can algorithms exhibit effort? They can (when implemented in computers) for sure play a lot of different sounds quickly and generate music of great complexity, but this does not correlate directly to effort, at least not on a human scale. These signs of human effort are, when exhibited by a computer, ineffective, because it lacks the markers of actual effort, which is an important part of human performance, such as the slight pause before a large intervallic jump on a string instrument or a large jump of the hands on a piano, because of the necessary preparation and the extra attention required from the performer.

It is hard to talk about machine effort when the effort is transferred from an effortless machine to our band-limited perception, which attempts cognitive parsing of the complexity. It becomes an effort in the listener instead of a perceived effort of the agent and completely lacks the empathetic dimension. It becomes just tiring to listen to. When listening to human-composed music, both these kinds of efforts may

be present. In today's AI-generated music, empathetic appreciation is still possible by mediation, to the humans behind it. This may change when AI becomes more autonomous.

As Arthur C. Clarke worded it in his third law: "Any sufficiently advanced technology is indistinguishable from magic." [14]. We have an attraction to this perceived magic, as we like to feel awe. And we who make music want to be the machine we do not understand. We like to not be completely understood, to be ambiguous or cause awe, and we seek the challenges in doing that.

31.6.4 Uncanny Valley

In robotics, they speak of the "uncanny valley", when a robot gets more and more human-like, suddenly it becomes scarily realistic, but you are not quite sure and something feels wrong [51]. Could this appear also in AI music? I do not think so, as we are already accustomed to so much of machine-generated and machine-mediated music-making. Also, technology has made new kinds of music possible, and acoustic instrumentalists have been influenced by this and developed new ways of playing that sound similar. For example, the acoustic drummer Jojo Mayer has reverse-engineered drum-and-bass style of playing, so today it is hard to tell what is machine-related or not, and our concept of music is extremely wide. If at all, I imagine it could appear related to form and perception, such as ill-proportioned distribution of musical ideas. An example could be music with sections of completely "wrong" durations, say, suddenly there are three hours of the same material in the middle of an otherwise perfectly composed song, a glitch that a human would be unlikely to produce. The large-scale form is still a challenge for generative music.

31.6.5 Authenticity

Could authenticity even be a thing, when no human agent is acknowledged? And what constitutes authenticity? It is related to honesty about how and by whom it was created (nominal authenticity), but also to empathetic appreciation and to effort. We can say it is authentic when it is created with an "honest effort" when the artist is doing "her best". Also, something is perceived as authentic when the artist is not doing it to please expectations, but to tell you something in earnest (expressive authenticity) [33].

But as we have seen, many AI algorithms are designed to give you the most probable outcome (to please, in a way), and as they have no intentional agency (yet), they are mere projections of influence from numerous human agents behind it at all design stages. And the question becomes one of *their* authenticity.

Is it possible that perceived authenticity would suffice, as when an artifact exhibits traces of process or of fictitious agents? Well, then it is based on a lie, and the goal of AI research supposedly is not to fake it, but to make it. It should be to

implement processes that can give us rich AI-generated art and music that could be aesthetically meaningful in the same way as human art and music.

In a future with autonomous AI, can AI music be authentic, and in what respect? It could be honest about where it comes from and embrace its mistakes, bugs, and limitations. It would be more honest and authentic if it were not masked as human art, and admit it is truly generated by AI. In the same way, it would be more authentic not to pretend to be AI-generated art, when humans are behind it, as is often the case today. So, current AI music and art is, in that sense, really not authentic, as it is presented as something it is not, in both directions.

31.6.6 Human Measure

All algorithmic and generative art and music (so far) has been made to be judged and received by humans. It may be machine generated, but it is made for human ears/eyes/brains—it is made to stay within the bandwidth of human perception, the frequency spectrum detectable by human hearing. When the human is taken out of the loop as sender and receiver, will the music lose its relevance to us? Will it become ungraspable?

Humans can only perceive phenomena at certain scales. Scientific instruments serve as translators to human measure and scale and to modes that we can perceive. They allow us to engage with phenomena of ungraspable scales, such as the large dimensions of space, or the microscopic scales of microorganisms and even atoms.

For whom is an AI presumed to make art? If the answer is humans, it becomes problematic. Where and how is the result coerced into human-perceivable form? And why? Does this not make the AI into an art-making slave to human masters?

If AI is supposed to make art and music for themselves, it becomes equally problematic. What complexity do machines need to achieve to need or want art? We can compare to animals, as being a little bit simpler beings than us, yet very complex—do they make or need art? We believe they do not, even though sexual selection has given rise to some spectacular aesthetic displays. That is, they are aesthetic to us, but we have no idea how they appear to the animals. And just as art produced by a different species could involve other frequencies of light or sound (many animals have senses different from ours), there is nothing that says that they would produce art perceivable to us. Instead, it is highly likely that it would be outside of our bandwidth, too complex or too simple, and incomprehensible to us. Why would it be related to our world at all, if produced by entities that live under completely different premises, who supposedly can communicate with each other through gigabit streams and which may have senses well beyond or totally different from our capabilities?

Their musical expressions may evolve into what to us sounds like super-complex noise or streams of thousands of notes per second. Or they may be completely uninterested in real-time streams and instead exchange large data objects to be parsed and analyzed in any customizable way. Or they may prefer digital silence as the optimal aesthetic experience. We can only speculate.

31.6.7 Cross-Species Art

If we assume autonomous artistic agents are possible in the future, with no (obvious) remaining connections to a human designer/engineer, or such connections being diluted by continued technical evolution or learning with no human intervention, their art and music will be for themselves, unless they explicitly make it for us because we ask them too. But would they even be interested in that? Or they make art for themselves, and we have to learn how to appreciate it, if at all possible—given that it is within our perceptual bandwidth and is communicable in media for which we have senses.

Would it be interesting then for them to reflect on our art? And is there any difference between these two directions of inter-species art? Would a cat be interested in our art or make art for us? Do we make art for cats?

Once general AI exists, it would take off in a path of its own, be completely disconnected from the human agency because of its complexity and capacity to learn and interact. It would most likely go beyond our abilities, not care about us. This is related to the extensive debate on the existential risks of AI (by Boström, Häggström and many others [11, 40, 55, 71]), but this is outside the scope of this text.

31.6.8 The Role of Time—Learning as a Non-Real-time Process

In AI (reinforcement) learning happens in simulated environments and can be computed in faster than real time (given enough computation power), and in parallel. Then, when it is applied in full in real-world action, this real world does not differ from the simulated worlds for the algorithms.

For humans, we have to live our learning experience. It has to happen in real time and not in parallel. We can isolate it somewhat into something called practice or rehearsal—I can practice my instrument or practice writing fugues and gradually get better at it. Then I apply my knowledge in a sharp situation, such as a concert or recording session.

There is a parallel here to composition versus improvisation. Composition is a non-linear, non-real-time process. I work on my composition until I am happy with it, or until I run out of time. In improvisation, the process unfolds in front of and in interaction with other musicians and other humans (listeners). If AI algorithms were forced to work at human time scales, their learning would be unbearably slow and they would not learn while we interact with them. The advantage of ML is still one of a brute-force batch process.

If the acceleration or compression of time that ML is capable of was applied to composition—as artistic creativity it is not a task that deals with virtuosity, but novelty and relevance of ideas in a given context—the compression still would not help. The extensive training may make it very good at imitating particular human composers (that it has been trained at), or perfect at some musical craft, but there is

nothing that says it would become better at coming up with a novel but relevant ideas.

There is a paradox hidden here. If we assume a composing algorithm could be trained, using reinforcement training, as if composition were a game, with quantifiable rewards (and sometimes the commercial musical life looks like that), then also the rewards would need to be estimated by a model of human listeners, or a model of human society and sociological mechanisms (as music popularity is as much sociology as it is aesthetics). But as this model needs to be trained in a real society or with real humans, since musical preference and musical style is a shifting thing, it cannot be accelerated without losing the connection to the real world. We would end up in the old dilemma of evolving/training composers and listeners/critics together, and there is a big risk that they will quickly diverge from human aesthetics and perceptual constraints.

31.6.9 Culture and Forgetting

We humans all start from scratch. We have to learn everything from other (typically older) humans, and also they had to learn from other humans. Tradition is a living thing that is constantly reinterpreted, and old ways are forgotten and replaced by new versions. It could be regarded as a process of continued refinement, getting more and more advanced, or better. Or as a sideways drift, growing different but not better.

An AI is thought of as constantly learning, yet, most ML systems converge toward a sufficiently good solution, and after a certain point, they don't learn anymore. If they over-learn, they will be too specific and lose the ability to generalize. But a human continues to develop.

We talk about different characteristic learning curves, e.g., the violin has a steep threshold to beginners, but the learning curve never ends, while the piano is easy to get started with, but certainly is not easy when you get to the advanced levels. For most acoustic instruments and most human skills in general, there is no clear end to the learning process. You can always get a little better. The same goes for composers. Maybe we cannot be better, but there is always something new to explore, and we get better at the meta-skills of managing our own development and at developing ideas.

We gradually change our ways and preferences, which also presumes a kind of forgetting as a condition for re-learning, which is believed to be beneficial for learning [45]. It may not be a loss of knowledge, but a gradual shift in the values and preferences that drives our actions. Such gradual forgetting is not handled gracefully in AI, with the well-known phenomenon of *catastrophic interference* as a good example: A system that learns something new will quickly and completely forget what it learned previously [47]. Some AI researchers think the key to forgetting will be a key to better AI [5].

31.7 Conclusions

31.7.1 Will AI Make Art-Making Easier?

Will musical skills no longer be needed, now and in the near future where generative AI algorithms can produce musical material for us?

I think it is still needed. One weak point of generative algorithms is large-scale form, and many projects involving AI-generated material still need humans to arrange the material, to do a working instrumentation. If you read the fine print, even projects like the AI-generated Beatles song mentioned earlier used human arranging and even human lyrics. As I see it, the lyrics, the arrangement, and the production (the “sound”) are very important parts of a pop composition, so it seems that the AI has a rather reduced influential agency in this case. Also, the human evaluating musical ear will always be needed, at least as long as we expect the algorithms to make music for humans.

There are also new skills that are needed to work with AI in music. You may need to learn how to construct sound engines and algorithms that are suitable for these techniques. Using AI algorithms effectively requires some understanding of how they actually work, like all tools. You may want to learn how to develop these kinds of tools as an extended meta-skill, especially if you want to take on the experimental artist/researcher role, wanting to go one step further.

So, there clearly is a craft of integrating AI into music, and it involves skills such as:

- Understanding complex systems.
- Understanding AI algorithms from both a theoretical and a practical point of view.
- Having an overview of available algorithms to be able to choose the right ones for the task.
- Understanding the potentials of specific algorithms and tools to be able to navigate their solution spaces.
- Understanding representations of music and how they constrain the output.
- Learning the new kinds of creative processes of generative tools (e.g., sow and harvest).
- Knowing the field to avoid doing the same things as others.

Richard Feynman famously wrote on his blackboard: “What I cannot create, I do not understand.”. He wanted to be able to derive something from the ground up, to understand all involved steps. This is maybe a utopian vision reserved for geniuses like him, but artists also need to understand their tools, practice on these algorithms to acquire sufficient skills, and form appropriate cognitive models and an intuitive understanding of what is possible, to be able to interact with them in musically meaningful ways. Artists should not rely on engineers to solve technical problems, because knowledge of the tools at hand shape our imagination, what we can think

of [24], and any delegation of the application of (or even more, the development of) tools involves (perhaps unknowingly) delegating aesthetic decisions.

Can anyone and everyone do this? I think they can. If the agency is proportionate to the inflow of information through interaction, everyone who is ready to put in enough effort can do it. It can be summarized as follows:

- **NO**,
there is no free lunch. With little effort spent and little input information flow, you can get fast but impersonal results, and your agency will be insignificant.
- **YES**,
with open-ended interactive generative systems, anybody with aesthetic judgment can spend enough time to breed/grow/generate characteristic and personal material, rendering interesting results. It may not be more efficient than other ways of working, but different. And your agency will be significant.

31.7.2 The Road Ahead—Musicking with Algorithms

Throughout this chapter, I have argued that most of today’s AI implementations, even though celebrated as autonomous creative agents in the popular press, are really just tools that mediate a distribution of human agency and that they generally lack creative agency of their own. The primary argument has been that many projects focus on the production of artifacts that are perceived as if they were created by a human artist, but this process is really one of mimetic optimization, which is inherently non-creative. We should apply AI algorithms to help us explore possibilities instead of optimizing toward known goals, or we will end up with conforming entertainment drones instead of AI artists. But we have to admit that modern AI algorithms are powerful tools of a new kind that operate at an abstraction level higher than tools previously available to artists and musicians. They do indeed bring amazing possibilities, but our expectations and agency attributions should be adequate. We must think critically about our tools, understand their inherent implications, but also learn them, use them, and develop the new crafts they deserve.

AI algorithms bring new modes of creation, such as the gardening paradigm, augmented creativity (e.g., interactive suggestion engines), and systems of creative human and machine agents. They offer fantastic new possibilities, but also new challenges and new crafts that nobody knows yet. If a composer wants to experience the fascinating possibilities of AI algorithms, she must let go of the urge for complete control, admit that algorithms bring something to the table, and see them as collaborators, interact with them, and share influential agency with them. This may mean not actually composing the notes of the piece, but instead taking a more curatorial role, tending to algorithms that generate new material or variations on her own material. If she spends significant time and effort with this kind of tool, she will undoubtedly still contribute significant agency and influence the result. Being aware

of these mechanisms may help a composer take the decision and work with them in a good way. And, she may choose what to delegate to the machine.

My generative AI systems will not slam the piano lid in frustration. They will not celebrate with a beer after a good gig, nor will they learn from a particularly humiliating situation on stage when things really did not work out. But they will indeed be a part of my process—I am happy to include them—and I need to understand what they can contribute, and how they can play with me.

Music-making is really an activity, it is *musicking* [64], a situated process full of interactions with information flowing in and out. We should embrace this process, as a human creative process and as implemented in algorithms. The aesthetic artifact is a by-product of the process. The interactions and their associated in- and outflow of information during the creative process fill the work with meaning—the interactions between artist and work, artist and environment, work and environment, algorithm and artist, and between algorithm and environment. The thoughts and computations that go into it, and the empathetic experience of the effort and struggle behind it, they both help to tell the story about its coming into being. We should stop our fixation with the artifacts and teach the AI algorithms the activity of musicking. We can include them in our musicking, and if and when they attain autonomy in the future, they may find their own way to music, instead of producing musical artifacts for humans.

AI algorithms can together with us be nodes in the network of creative agents that together with us shape the art. We integrate with them in an activity of situated creative intelligence. For an even stronger kind of shared interactive creativity, we may turn our interest toward smaller and simpler AI models that can be used interactively and integrate with musicians in real-time embodied interaction, i.e., we are playing not *on* but *with* them. Again, together with them, we become part of a shared situated creative musicking intelligence.

As long as we enforce human measures onto AI art, it will be us who create it, as we impose serious constraints and expectations onto it. While autonomous creative machines are theoretically possible, they are far, far away. But if and when they come, they will probably not make art for us. While waiting, let us play with them.

References

1. Anscombe, G. E. M. (1957). *Intention*. Oxford: Basil Blackwell.
2. Ariza, C. (2009). The interrogator as critic: The turing test and the evaluation of generative music systems. *Computer Music Journal*, 33(2), 48–70
3. Ashby, W. R. (1957). *An Introduction to Cybernetics*. Chapman and Hall.
4. Barthes, R. (1977). The Death of the Author, in *IMAGE-MUSIC-TEXT, Essays selected and translated by Stephen Heath*, pp. 142–148. Hill and Wang.
5. Beierle, C., & Timm, I. J. (2019). Intentional forgetting: An emerging field in AI and beyond. *KI - Künstliche Intelligenz*, 33, 5–8.
6. Belgum, E., Roads, C., Chadabe, J., Tobenfeld, T.E., & Spiegel, L. (1988). A turing test for “musical intelligence”? *Computer Music Journal*, 12(4), 7–9.

7. Berrar, D. P., & Schuster, A. (2014). Computing machinery and creativity: Lessons learned from the turing test. *Kybernetes*.
8. Bialek, W., Nemenman, I., & Tishby, N. (2001). Predictability, complexity, and learning. *Neural Computation*, 13(11), 2409–2463
9. Biles, J. A. (1994). GenJam: A genetic algorithm for generating jazz solos. In *Proceedings of the International Computer Music Conference, Aarhus, Denmark*, pp. 131–137.
10. Boden, M. A. (2010). The turing test and artistic creativity. *Kybernetes*.
11. Boström, N. (2014). Superintelligence: Paths, dangers, strategies. Oxford University Press.
12. Bown, O. (2015). Attributing creative agency: Are we doing it right? In *Proceedings of the International Conference on Computational Creativity (ICCC)*, pp. 17–22.
13. Bown, O., & McCormack, J. (2009). Creative agency: A clearer goal for artificial life in the arts. In *European Conference on Artificial Life*, pp. 254–261. Springer.
14. Clarke, A. C. (1973). *Profiles of the future; an inquiry into the limits of the possible* (rev. ed.). Harper and Row.
15. Cope, D. (2005). *Computer Models of Musical Creativity*. Cambridge, Mass.: MIT Press.
16. Cross, L., & Tolbert, E. (2009). Music and meaning. In *The Oxford Handbook of Music Psychology*, 2nd ed. Oxford University Press.
17. Sony, C. S. L. (2020). Daddy’s car: A song composed by artificial intelligence—in the style of the beatles. Retrieved from July 27, 2020, from https://youtu.be/LSHZ_b05W7o.
18. Dahlstedt, P. (2001). Creating and exploring huge parameter spaces: Interactive evolution as a tool for sound generation. In A. Schloss, R. Dannenberg, & P. Driessen (Eds.), *Proceedings of the International Computer Music Conference*. Habana, Cuba: Instituto Cubano de la Musica, and San Francisco, CA: International Computer Music Association, pp. 235–242.
19. Dahlstedt, P. (2004). *Sounds Unheard of: Evolutionary algorithms as creative tools for the contemporary composer*. Ph.D. thesis, Chalmers University of Technology.
20. Dahlstedt, P. (2007). Autonomous evolution of complete piano pieces and performances. In F. A. e Costa et al. (Eds.), *Proceedings of MusicAL Workshop, Advances in Artificial Life, 9th European Conference, ECAL 2007, Lisbon, Portugal, September 10–14, 2007, Proceedings (Workshop CDROM)*. Lecture Notes in Computer Science 4648 Springer.
21. Dahlstedt, P. (2007). Evolution in creative sound design. In E. R. Miranda & J. A. Biles (Eds.), *Evolutionary Computer Music*, pp. 79–99. London: Springer.
22. Dahlstedt, P. (2009). Thoughts on creative evolution: A meta-generative approach to composition. *Contemporary Music Review*, 28(1), 43–55
23. Dahlstedt, P. (2011). Autonomous evolution of piano pieces and performances with ossia ii. In E. R. Miranda (Ed.) *A-Life for Music: Music and Computer Models of Living Systems*. A-R Editions.
24. Dahlstedt, P. (2012). Between material and ideas: A process-based spatial model of artistic creativity. In J. McCormack, & M. d’Inverno, (Eds.), *Computers and Creativity*, pp. 205–233. Springer.
25. Dahlstedt, P. (2018). Action and perception: Embodying algorithms and the extended mind. In R. T. Dean, & A. McLean (Eds.) *OUP Handbook of Algorithmic Music*, Chap. 3, pp. 41–65. Oxford University Press.
26. Dahlstedt, P. (2019). Big data and creativity. *European Review*, 27(3), 411–439
27. Dahlstedt, P. (2020). Ossia suite (musical work). Retrieved July 27, 2020, from <https://soundcloud.com/palle-dahlstedt/sets/ossia-suite>.
28. Dahlstedt, P. (2001). A mutasynth in parameter space: Interactive composition through evolution. *Organised Sound*, 6, 121–124.
29. Dahlstedt, P. (2015). Turn-based evolution in a simplified model of artistic creative process. *Evolutionary Intelligence*, 8(1), 37–50
30. Dahlstedt, P., & McBurney, P. (2007). Musical agents: Toward computer-aided music composition using autonomous software agents. *Leonardo*, 39(5), 469–470.
31. Dahlstedt, P., & Nordahl, M. G. (2001). Living melodies—Coevolution of sonic communication. *Leonardo*, 34(3), 243–248

32. Davidson, D. (1963). Actions, reasons, and causes. *The Journal of Philosophy*, 60(23), 685–700
33. Dutton, D. (2003). Authenticity in art. In *The Oxford Handbook of Aesthetics*, pp. 258–274.
34. Garrabrant, S., & Demski, A. (2018). Embedded agency (full-text version). Alignment Forum. Retrieved July 27, 2020, from <https://www.alignmentforum.org/posts/i3BTagvt3HbPMx6PN/embedded-agency-full-text-version>.
35. Gregory, R. L. (1981). *Mind in Science*. London: Weidenfeld and Nicolson.
36. Hewitt, C. (1977). Viewing control structures as patterns of passing messages. *Artificial Intelligence*, 8(3), 323–364
37. Hewitt, C., Bishop, P., & Steiger, R. (1973). A universal modular actor formalism for artificial intelligence. In *Proceedings of the 3rd international joint conference on Artificial intelligence (IJCAI)*, pp. 235–245.
38. Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.
39. Huron, D. B. (2006). *Sweet anticipation: Music and the psychology of expectation*. MIT press.
40. Häggström, O. (2016). *Here be dragons: Science, technology and the future of humanity*. Oxford University Press.
41. Johnson, D. G., & Verdicchio, M. (2019). AI, agency and responsibility: The VW fraud case and beyond. *AI & SOCIETY*, 34(3), 639–647.
42. Knuth, D. E. (1968). *The Art of Computer Programming*, Vol. 1: Fundamental Algorithms. Addison-Wesley.
43. Latour, B. (2005). *Reassembling the Social: An Introduction to Actor-Network Theory*. Oxford: Oxford University Press.
44. Lehman, J., Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pp. 329–336.
45. Markovitch, S., & Scott, P. D. (1988). The role of forgetting in learning. In *Machine Learning Proceedings 1988*, pp. 459–465. Elsevier.
46. McCarthy, J. (2020). Sony uses ai to compose a beatles-inspired song, the unforgettable 'daddy's car'. Retrieved July 27, 2020, from <https://www.thedrum.com/news/2016/09/27/sony-uses-ai-compose-beatles-inspired-song-the-unforgettable-daddy-s-car>.
47. McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of Learning and Motivation*, vol. 24, pp. 109–165. Elsevier.
48. McCormack, J., & Dorin, A. (2001). Art, emergence and the computational sublime. In *Proceedings of the Second International Conference on Generative Systems in the Electronic Arts, Victoria, Australia, Centre for Electronic Media Art, VIC, Australia*.
49. Meyer, L. B. (1957). Meaning in music and information theory. *The Journal of Aesthetics and Art Criticism*, 15(4), 412–424.
50. Minsky, M. (1982). Music, mind, and meaning. In *Music, Mind, and Brain*, pp. 1–19. Springer.
51. Mori, M. (2012). The Uncanny Valley [From the Field]. *IEEE Robotics & Automation Magazine*, 19(2), 98–100.
52. Mike Murphy, M. (2020). Computers can now paint like van gogh and picasso. Retrieved July 27, 2020, from <https://qz.com/495614/computers-can-now-paint-like-van-gogh-and-picasso/>.
53. Nielsen, P. (1972). Some comments on Vagn Holmboe's idea of metamorphosis. *Dansk årbog for musikforskning*, pp. 1968–1972.
54. Nilsson, P. A. (2011). *A Field of Possibilities: Designing and Playing Digital Musical Instruments*. Ph.D. thesis, University of Gothenburg.
55. Omohundro, S. (2014). Autonomous technology and the greater human good. *Journal of Experimental & Theoretical Artificial Intelligence*, 26(3), 303–315

56. Pachet, F. (2012). Musical virtuosity and creativity. In *Computers and Creativity*, pp. 115–146. Springer.
57. Pease, A., & Colton, S. (2011). On impact and evaluation in computational creativity: A discussion of the turing test and an alternative proposal. In *Proceedings of the AISB symposium on AI and Philosophy*, vol. 39.
58. Russell, S. J., & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach* (2nd ed.). Upper Saddle River, New Jersey: Prentice Hall.
59. Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., & Müller, K. R. (2019). *Explainable AI: Interpreting, explaining and visualizing deep learning*. Vol. 11700. Springer Nature.
60. Sayes, E. (2014). Actor–network theory and methodology: Just what does it mean to say that nonhumans have agency? *Social Studies of Science*, 44(1), 134–149
61. Schlosser, M. (2019). Agency. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2019 edition.
62. Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(379–423), 623–656
63. Sims, K. (1991). Artificial evolution for computer graphics. In *ACM SIGGRAPH '91 Conference Proceedings, Las Vegas, Nevada*, pp. 319–328.
64. Small, C. (1998). *Musicking: The meanings of performing and listening*. Wesleyan University Press.
65. Tawaststjerna, E. (1976). *Sibelius vol. 1*, translated by R. Layton. London: Faber and Faber.
66. Tishby, N., & Polani, D. (2011). Information theory of decisions and actions. In *Perception-action Cycle*, pp. 601–636. Springer.
67. Touchette, H., & Lloyd, S. (2004). Information-theoretic approach to the study of control systems. *Physica A: Statistical Mechanics and Its Applications*, 331(1–2), 140–172
68. Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.
69. Vygotsky, L. S. (1978). *Mind in Society*. Cambridge, MA: Harvard University Press.
70. Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press.
71. Yampolskiy, R. V. (2014). Utility function security in artificially intelligent agents. *Journal of Experimental & Theoretical Artificial Intelligence*, 26(3), 373–389.

Palle Dahlstedt Palle Dahlstedt studied piano, composition, and electronic music (M.F.A. and M. A.), and has a Ph.D. in the topic of evolutionary computation for artistic creativity from Chalmers University of Technology, Sweden. His research focuses on new technologies for improvisation, composition, and art. He is particularly interested in advanced algorithms for creativity, technologies that enable embodied performance with electronic sounds, and new kinds of interactions based on a systems view of emergence from human–technology interactions. Palle is an associate professor of Interaction Design at the Department of Computer Science and Engineering of the University of Gothenburg and teaches composition at the Academy of Music and Drama, Gothenburg, Sweden. He also teaches at Chalmers University of Technology in Sweden and is a guest professor at Aalborg University, Denmark. E-mail: palle@chalmers.se