

# From Artificial to Extended Intelligence in Music Composition

ARTEMI-MARIA GIOTI

Institute of Electronic Music and Acoustics (IEM), Graz, Austria  
Email: [gioti@iem.at](mailto:gioti@iem.at)

**This article explores the relationship and disparities between human and computational creativity by addressing the following questions: How well are computational creativity systems currently performing at creative tasks? Could computers outperform human composers? And, if not, is computational creativity a utopia? Automatic composition systems are examined with respect to Boden's three criteria of creativity (*novelty*, *surprise* and *value*), as well as their assumptions about the nature of creativity. As an alternative to a competitive relationship between human and computational creativity, the article proposes the concept of a distributed human–computer co-creativity, in which computational creativity extends – rather than replaces – human creativity, by expanding the space of creative possibilities.**

## 1. INTRODUCTION

Boden (2004: 1) defines creativity as ‘the ability to come up with ideas or artefacts that are new, surprising and valuable’. She identifies two ways in which an idea (or artefact) can be *new* and consequently distinguishes between two types of creativity: psychological and historical (*P-creativity* and *H-creativity* respectively). *P-creativity* involves an idea that is new to the person who conceived it, while *H-creativity* involves an idea that is historically new, that is, an idea that has been conceived for the first time in human history. Similarly, there are three ways in which an idea can be *surprising*: it can be unlikely, unexpected but fitting into an existing *conceptual space* (i.e., style of thought) or previously thought of as impossible. While Boden's first two criteria (novelty and surprise) are relatively straightforward and unambiguous, the third one (value) resists a precise definition, since, as Boden points out, aesthetic values are not only difficult to describe, but also vary across cultures, or even subcultures within the same culture; and of course they change through time (Boden 2010: 39).

Dorin and Korb (2012) propose an alternative definition of creativity that focuses exclusively on novelty, rejecting notions of value and appropriateness as irrelevant. They specifically criticise discussions of value as counterproductive for computational creativity and non-essential to understanding human creativity.

They suggest that what makes an activity creative must relate to the activity itself rather than the reception of its outcomes. Bown (2012), on the other hand, distinguishes between two types of creativity: *generative* and *adaptive*, the difference between the two being that only the latter is concerned with value. This debate on value will be examined more closely later in this article. For now, Boden's definition will be used as a guide in determining to what extent computer-generated artefacts fulfil her criteria of creativity and what implicit and explicit assumptions about human creativity are evident in the design of computational creativity systems.

While creativity can be understood in a variety of contexts, spanning over a broad spectrum of human and non-human activity – creativity can be attributed not only to humans, but also to biological systems and processes, such as evolution (Bentley and Corne 2002) – this article will focus on creativity as it relates to artistic production. Specifically, it will focus on art music composition, including both acoustic and electroacoustic music composition, and try to answer the following questions: How well are computers currently performing at creative tasks? Could computers outperform human composers? And, if not, is computational creativity a utopia?

## 2. HOW WELL ARE COMPUTERS PERFORMING AT CREATIVE TASKS?

Automatic composition and music generation systems cover a wide range of musical styles and genres and therefore employ different types of data and algorithms. While an exhaustive literature review of such systems is beyond the scope of this article, the three examples discussed below are meant to illustrate different approaches to the automation of musical tasks, along with their shortcomings, as reported by their designers and other researchers.

David Cope's *EMI* (Experiments in Musical Intelligence) is one of the most well-known automatic composition systems designed for acoustic composition. *EMI* performs statistical analysis on a corpus of sample works stored as MIDI scores and identifies

patterns found in more than one of the sample works. It then uses Augmented Transition Networks (ATNs) in order to generate new works. The generated compositions are finally analysed and compared to the sample works. Cope's system focuses exclusively on pitch and duration as a means to reduce the dimensionality of the input data and does not take into account parameters such as timbre and dynamics. Cope identifies as a weakness of the program its bias towards 'diatonic tonal music', suggesting that it could be expanded to recognise the minor mode and other scales (Cope 1992: 82). He also suggests that the system could benefit from additional software components written to address chromaticism, cadences, phrase length and musical form.

Another automatic composition system, *Autocousmatic*, generates electroacoustic art music compositions using a database of sound files provided by the user. The input sound files are processed and machine listening is used in order to avoid 'overloads and other digital nastiness, as well as silence and low activity' (Collins 2012: 10). The form of the generated mixes is built based on a section duration and density model, derived through manual analysis of sample works. An optional component can be used to evaluate the generated output mixes with respect to their proximity to Trevor Wishart's *Vox 5*. The similarity is determined using Dynamic Time Warping (DTW) and features such as perceptual loudness, sensory dissonance, onsets and several spectral descriptors (Collins 2012). As part of the system's evaluation, *Autocousmatic*-generated compositions were submitted to music festivals and conferences and feedback was sought from professional electroacoustic music composers. Collins reports that none of the submissions has been successful so far. The professional composers asked to evaluate the system expressed criticism towards its ability to generate larger forms, describing the transitions between different sections as 'abrupt' and 'arbitrary' and criticising the generated mixes for lacking 'directionality' (ibid.).

A more recent trend in automatic music generation involves the use of Deep Learning algorithms that learn from unstructured data (i.e., raw audio), such as *WaveNet* (van den Oord et al. 2016). *WaveNet* was developed mainly for speech applications and is based on a probabilistic and autoregressive model, that is, a model in which predictions for each audio sample are conditioned on all previous samples (van den Oord et al. 2016). *WaveNet* can generate musical sequences (in the form of raw audio) with partially convincing local structure, but poor global structure, showing that the algorithm fails to learn mid- and long-range dependencies (Manzelli, Thakkar, Siahkamari and Kulis 2018). Furthermore, music applications of *WaveNet* (van den Oord et al. 2016; Manzelli et al. 2018) have so far focused on tonal, pitch-based music, in which

short- and long-term dependencies are governed by the rules of traditional harmony. Whether the algorithm would perform better or worse on a corpus of electroacoustic music or even contemporary instrumental music remains to be seen.

### 3. COULD COMPUTERS OUTPERFORM HUMANS IN CREATIVE TASKS?

Creativity involves highly complex cognitive and psychological processes, a simulation of which would be an undoubtedly ambitious undertaking. As a way to reduce data dimensionality and computational complexity, automatic composition systems employ models based entirely on domain-specific features as input data, such as MIDI data (Cope 1992) or perceptual audio descriptors (Collins 2012), and reframe the problem of musical creativity as one of style imitation. This is evident in the use of musical corpora and techniques such as statistical analysis and machine learning in order to model and subsequently produce artefacts within certain musical styles. An obvious limitation of this approach lies in its prioritisation of mastery over innovation. Presumably, even if a style imitation system succeeded in 'mastering' the style of one or more human composers (i.e., if it passed the Turing test), it still would not be able to produce anything innovative, as that would be beyond its scope and intent.

Creativity, particularly *H-creativity*, is not just a matter of mastery, but also imagination, resourcefulness and invention. In fact, some of the most pivotal works in music history are those that broke away from tradition, either by proposing new composition systems (e.g., the twelve-tone system), or through their innovative and imaginative use of technology (e.g., Steve Reich's *Pendulum Music* (1968)), or even by questioning the very definition of music (e.g., John Cage's *4'33"* (1952)). For that reason, even if computers achieved a high level of mastery within certain styles, we would still need to investigate whether they are capable of creating new styles, an objective that would require fundamental changes to the models currently employed by computational creativity systems.

Most subfields of Artificial Intelligence (AI), from rule-based systems to Machine Learning (ML), share a common understanding of artificial intelligence as the ability of computers to solve domain-specific problems commonly associated with humans. Machine learning brought about a revolution in the way these tasks were performed: instead of executing a set of hand-coded rules, computers could now learn by examples, figuring out the rules 'on-the-fly'. This made possible many applications that would be nearly impossible otherwise (e.g., image and speech recognition). However, when it comes to the automation of musical tasks, both rule-based systems and machine

learning applications seem to be limited in their definition of *creativity*. For example, both training examples and rules are derived from already existing artefacts. They are manifestations of the stylistic constraints of already existing styles. Creativity is then understood as the ability to imitate, or conform to the constraints of a given style, as discussed above.

Besides *P-creativity* and *H-creativity*, Boden makes an additional distinction between *combinational*, *exploratory* and *transformational* creativity (Boden 2004). *Combinational creativity* involves combining familiar ideas in unfamiliar and surprising ways. *Exploratory creativity* involves the generation of new ideas within an existing conceptual space (e.g., an existing style). Finally, *transformational creativity* involves transforming and essentially redefining an existing conceptual space (i.e., creating new styles). Interestingly, Boden mentions Schoenberg's atonality as an example of transformational creativity (Boden 2007), while she considers automatic composition systems, such as David Cope's *EMI*, as examples of exploratory creativity (Boden 2004).

Assuming that automatic composition systems are indeed capable of exploratory creativity, then the question that needs to be answered is: are computers capable of transformational creativity? Boden answers this question positively and uses evolutionary algorithms as an example of how computer programs can randomly change their rules, thereby transforming their conceptual space (Boden 2010: 38).

However, in assuming that by randomly generating a new rule system the computer has created a new 'style', Boden has overlooked her third criterion of creativity: *value*. Whether a deviation from existing styles constitutes an anomaly or a paradigm shift is not only a question of novelty, but also impact. A deviation from the norm alone does not qualify as a new style if it is historically inconsequential. Whether a new, human- or computer-generated rule system qualifies as a 'style' can only be determined by its acceptance – or lack thereof – by a society or social group (de Jager 1972), such as contemporary composers, and/or its replication by other members within that group (Meyer 1983). Schoenberg's twelve-tone system is considered a 'style' because of its impact on Schoenberg's contemporaries and successors. Had it not had such an impact on music history, the 12-tone system would probably not hold the same cultural value it does today.

Dorin and Korb (2012) attempt to detach creativity from notions of value, suggesting that what makes an activity creative must relate to the activity itself, rather than the reception of its outcomes. They use the example of artists who were recognised posthumously to illustrate that value relates to reception, not creativity. However, even if we accept that what makes an

activity creative is determined by the activity itself, we still have to acknowledge that what makes the outcome of this activity *art* is determined by reception. Marcel Duchamp's *Fountain*, a readymade sculpture consisting of a urinal, is only art because we (society, or a smaller group within it) agree it is art. The cultural value attributed to the object is a result of extrinsic, not intrinsic qualities: the urinal itself is a mass-produced commercial product. It seems then that, at least in the case of art, creativity is inextricably connected to value. As for cases of posthumous recognition, one might argue that these are proof of the dynamic and complex nature of value systems, not their irrelevance. Most importantly, evaluation is carried out not only at reception, but also as part of the creative process. During the latter, ideas are constantly evaluated by the artist based on their subjective, but culturally informed, values.

Duchamp's example illustrates that creativity, especially transformational creativity, cannot be reduced to a set of domain-specific skills, nor studied outside a broader social context. Creativity is a situated phenomenon driven among others by social, cultural, psychological, political, technological and economical factors. Nowhere is this better illustrated than in the work of pioneers such as Pauline Oliveros and John Cage.

Pauline Oliveros's *Sonic Meditations*, a collection of verbal scores for a group of performers/participants meeting regularly over a longer period of time, is a revolutionary work due, among other reasons, to its participatory approach to music-making. Musical training is not a requirement for participation, since, as Oliveros states, she aims to 'erase the subject/object or performer/audience relationship by returning to ancient forms which preclude spectators' (Oliveros 1974). Expanding one's sonic awareness, sharing a common experience with other members of the group and releasing psychological and physiological tension are also mentioned as some of the goals of this activity, of which 'music is a welcome by-product' (Oliveros 1974). This last sentence is enlightening with regard to the composer's aims and priorities: Oliveros is more interested in the social, psychological and even physiological aspects of music-making than in its product.

Similarly, in *4'33"* John Cage (1952) poses a series of ontological questions ('What is music?', 'What constitutes a musical work?', etc.) in a piece that consists entirely of silence. The aesthetic and cultural value of the piece, as well as its novelty, lies in the position it takes with respect to the debate on musical ontology – another position being Varèse's definition of music as 'organised sound' (Varèse and Wen-Chung 1966). In this particular example, taking the work out of its context would be stripping it of its cultural value. For example, one could not expect to analyse this

piece, feed the data into a machine learning algorithm and generate music in Cage's style. While this might seem as an extreme example, it illustrates in a practical way that creativity in art cannot be understood outside a sociocultural context.

Another aspect of art creation and reception exemplified by Cage's work is that of inter-human communication. From the composer's perspective, *4'33"* is based on certain assumptions about listeners' general and even specialised knowledge, as well as their expectations of a musical work. From a listener's perspective, reception and interpretation are based on similar assumptions (e.g., that the composer is human and is communicating some thoughts, albeit through unconventional means). Even when expectations are subverted, this is interpreted as an intentional act of communication. Having said that, if we accept O'Hear's (1995) definition of art as inter-human communication, then we axiomatically reject computer-generated art as an impossibility.

This article wishes to adopt a pragmatic approach and will therefore avoid shifting the discussion into philosophical debates on consciousness and intentionality. The examples mentioned above are just meant to illustrate an obvious shortcoming in current approaches to the automation of compositional tasks: the assumption that creativity is a domain-specific skill and that the domain at hand (music) can be studied in isolation from any social context.

Admittedly, these examples involve a specific type of innovation: *innovation in goals* (de Jager 1972). *Innovation in goals* refers to innovation with respect to extra-musical ends (e.g., activist art), while *innovation in means* can refer among others to the invention of new composition systems (e.g., serialism), new instruments (e.g., sensor-based interfaces), or new performance practices (e.g., live coding). Particularly in contemporary music creation, however, it can be rather hard to distinguish between the two types of innovation, since in many cases they seem to co-exist. For example, participatory art can be considered as innovative with respect to both its ends (lifting the distinction between performer/author and spectator) and its means (technologies and practices that enable and encourage participation). Oliveros's *Sonic Meditations* is a good example of a concurrence of the two types of innovation, in which a novel artistic goal (exploring the social aspects of music-making through participation and inclusion) leads to the invention of new means, designed to accommodate that new goal (verbal scores that can be interpreted by non-musicians).

However, even in Schoenberg's case, in which innovation is restricted to the means, transformational creativity required an extensive knowledge of music history, leading to the realisation that the tonal system had reached and exceeded its limits. This illustrates that, even when innovation is restricted to

the means, the *why* still matters and is subject to a cultural and historical context.

Additionally, innovation in means can be driven or influenced by extra-musical factors, such as technological and economical advances. For example, *musique concrète* would not have existed if it were not for recording technology. Similarly, a laptop orchestra would have been unthinkable before the invention, in 1966, of the integrated circuit (microchip) and that of the personal computer less than two decades later. The small size, low-cost and high-computational power of modern computers have revolutionised the way in which music is composed and performed. Network performances, interactive human-computer improvisation and performances with sensor-based interfaces are only a few examples of artistic practices facilitated by technological and economical advances.

The task of defining the factors that can play a role in creative decision-making becomes exponentially more complex when considering what is probably the most essential component of human creativity: psychology. A person's subjective experience of the world is perhaps the least quantifiable of all factors influencing creative decision-making. Real-life experiences and the cognitive and emotional responses they may trigger in an individual differ from one person to another and are nearly impossible to simulate and predict. Artists, unlike software agents, are *subjects* with unique personalities, aesthetic preferences and belief systems (values, opinions, etc.), all of which influence creative decisions.

By overlooking extra-musical factors that influence creative decision-making, automatic composition systems are making either one of the following implicit or explicit assumptions: 1) that creativity in music is independent of any social context and can be modelled using domain-specific features; or more likely, 2) that modelling transformational creativity as a social, situated phenomenon is currently beyond computational means and therefore all computational creativity can aim for is the exploration of existing musical styles.

Interestingly, specifically in acoustic music composition, computational creativity systems have focused on musical styles governed by the rules of tonal harmony. For instance, there currently appears to be no automatic composition system dealing with contemporary art music composition. In the latter, 'style' cannot be defined in terms of pitch or harmony, while there is no universally agreed-upon notation system, with each composer essentially creating and using their own. A MIDI reduction of musical scores of this type would be practically useless – if at all possible.

To formulate the problem in machine learning terms, the objective of automating compositional tasks seems to pose a significant challenge for feature engineering. Describing musical styles in terms of stylistic constraints, using domain-specific features (e.g., MIDI data

or audio descriptors) might be adequate for the purpose of simulating exploratory creativity – at least within certain styles. But, for computational creativity to go beyond exploratory creativity (i.e., style imitation), computational models would have to be adapted accordingly to reflect the *situatedness* of human creativity, taking into account extra-musical (sociopolitical, cultural, psychological, technological, etc.) factors that influence creative decisions.

A counter-position to this argument might be that computational creativity does not need to simulate human creativity. This may apply particularly to cases of human–computer co-creativity, in which computational creativity is understood as complementary to human creativity. However, the question being asked here is whether computational creativity systems, specifically autonomously creative systems, are capable of transformational creativity. The work of pioneers such as Cage, Oliveros, Schoenberg and many others suggests that transformational creativity requires knowledge and interpretation of a sociocultural and historical context. It should follow then that, in order for computers to be capable of transformational creativity, they should have similar capabilities.

While this remains beyond computational means, computational creativity will not be able to challenge human creativity. That is not to say that computers will never be capable of transformational creativity but rather that, if we were to pursue this objective, we would have to devise models that reflect the situatedness of human creativity and avoid oversimplified assumptions that equate creativity to style imitation.

At this point a clarification is needed: the claim that while computational creativity remains limited to exploratory creativity it will not be able to challenge human creativity refers to the aesthetic and cultural value of computer-generated art, not its financial value. In fact, computational creativity is already producing high financial value and will probably continue to do so. For example, in October 2018, an eighteenth-century-style painting generated by a generative adversarial network was sold for over \$400,000 at a well-known auction house (Cohn 2018). However, market value is not to be confused with cultural or aesthetic value, nor is it an indicator of innovation. Whether this AI-generated painting will have an influence on the artistic community remains to be seen. Regarding its potential for innovation, that is, as expected, rather low, since it was generated by a model trained on already existing paintings. One could, of course, argue that innovation here lies in the technology through which it was produced, not the artefact itself. Even so, the painting remains an example of exploratory creativity, as technological innovation does not constitute aesthetic innovation.

#### 4. EVALUATION OF COMPUTER-GENERATED ARTEFACTS

Besides feature engineering, another significant challenge for computational creativity is the evaluation of computer-generated artefacts. The evaluation of computer-generated compositions is usually based on Turing-like tests designed to determine whether they are distinguishable from compositions created by humans. A few variations of the Turing Test have been proposed specifically for the evaluation of generative music systems (Ariza 2009). While the rationale behind this approach is understandable, it is still important to note that a Turing Test says little about the aesthetic value of an artwork: musical works are not normally assessed with respect to their believability. Furthermore, evaluation in composition is formative, rather than summative. This means that evaluation takes place not only after a composition is completed, but also while it is in progress. The integration of machine listening processes in automatic composition systems such as *Autocousmatic* (Collins 2012) is a first step towards addressing the issue of formative evaluation in the automation of musical tasks. However, it is still worth pointing out that listening as an act of aesthetic appreciation, based on culturally informed and most importantly *subjective* aesthetic criteria, is still far from being simulated by Music Information Retrieval.

In addition to Turing Tests, computer-generated artefacts can also be evaluated through computational means. The history of computational aesthetic evaluation encompasses a wide range of methods and approaches, from formulaic theories and biologically inspired fitness measures to psychological models and empirical studies of human aesthetics (Galanter 2012). Galanter (2012) distinguishes between two modes of computational aesthetic evaluation based on whether the aesthetic standards are defined by humans or generated by software agents, but goes on to point out that software-generated aesthetics often ‘feel alien and disconnected from human experience’ (Galanter 2012: 256). He claims that research in psychology and neurology has shown promising results for future work in computational aesthetic evaluation, but acknowledges its limitations with regard to the highly subjective and complex nature of human aesthetic evaluation.

McCormack is more critical towards universal aesthetic values derived through empirical studies (e.g., Martindale 1999), arguing that aesthetic preferences depend on cultural values and individual taste and that ‘surface aesthetic qualities’ often have little relevance for the appreciation of contemporary art (McCormack 2012: 44).

The idea of removing the artist’s aesthetics from the creative process and replacing it through

computer-generated aesthetics has both been praised (Dorin and Korb 2012) and met with scepticism (McCormack 2012; Galanter 2012). All in all, the debate on the evaluation of computer-generated artefacts is ongoing and far from being settled.

## 5. IS COMPUTATIONAL CREATIVITY A UTOPIA?

Given the difficulties involved in both the automation of creative processes and the evaluation of computer-generated artefacts, one might ask whether computational creativity is at all possible. The answer to that question strongly depends on our definition and expectations of computational creativity. Expecting that computers will generate historically consequential art of high cultural value is a rather utopian vision, or dystopian, if we consider the ethical and economic impact of computers displacing human artists. However, for computational creativity to be worth pursuing it does not have to compete with human creativity.

The question of whether human and computational intelligence should be in a competitive or complementary relationship with each other recalls the debate on *Artificial Intelligence* (machines replacing humans at cognitive tasks) versus *Intelligence Augmentation* (machines assisting humans at cognitive tasks) (Licklider 1960; Engelbart 1962; Ashby 1964). AI has led to impressive results in controlled environments, where the inputs and goals of the algorithm are clearly defined and its performance can be evaluated in quantitative terms; for example, in applications such as image and speech recognition, or even self-driving cars. However, in creative tasks, where the inputs and goals are hard to define, AI has demonstrated less promising results. This suggests that, when it comes to creative tasks, we might need to think of artificial intelligence in different terms.

Ito proposes the term *Extended Intelligence* (EI), instead of Artificial Intelligence, to indicate an understanding of intelligence as a ‘fundamentally distributed phenomenon’. He argues that AI should be seen as yet another actor contributing to a ‘networked intelligence’ that encompasses both humans and machines (Ito 2016). Concepts such as collective intelligence and collective learning (the process through which knowledge and information is shared and preserved across generations in human societies) could have a transformative potential for the way we conceptualise computational intelligence. Collective learning implies that intelligence extends beyond the individual mind to a network – a society – of minds. Computational intelligence could then contribute to a networked intelligence by augmenting, not replacing, human intelligence.

A connection between the idea of an extended intelligence contributing to a collective human–machine intelligence and Latour’s Actor-Network-Theory is

easy to draw. Latour defines an actor as an entity which ‘is made to act’ (Latour 2005: 46) and as ‘any thing’ that modifies a state of affairs (Latour 2005: 71), pointing out that objects, not just humans, can be actors. Latour does not fail to acknowledge the asymmetry between human and non-human actors: human and non-human actors do not have the same type of agency, a distinction that is crucial in regarding extended intelligence as an actor. Extended intelligence undeniably has the potential to modify the ‘state of affairs’ in music composition, despite its relationship to human actors being asymmetrical, and therefore falls under Latour’s definition of an actor.

## 6. WHAT ARE THE IMPLICATIONS OF EXTENDED INTELLIGENCE FOR MUSIC COMPOSITION?

The idea of an extended and distributed intelligence in music composition translates into practices that encompass both human and machine actors. This, in practice, means shifting the focus from automating creative tasks to re-conceptualising those tasks with the help of machine intelligence. The purpose of such an approach is to potentially extend what is creatively possible and gain a better understanding of human creativity as a whole.

In line with that approach, McCormack’s concept of *creative ecosystems* encompasses ‘humans, technology and the socially/technologically mediated environment’. His approach to creativity views the creative process as an explorative, rather than an optimisation process, in which human creativity is enhanced through computational means. The ecosystemic approach does not aim to automate the creative process or replace human aesthetic judgement, but rather open up new possibilities for creative exploration, allowing the artist to expand their creativity (McCormack 2012).

In music composition, applications of such a distributed human–computer co-creativity can be found among others in computer-assisted composition and interactive music systems, that is, systems in which human performers interact with software agents in a reciprocal manner. In interactive music systems machine intelligence extends what is humanly possible (e.g., through generative rule-based processes and fast information processing), while high-level aesthetic decisions are made by the human. Jones, Brown and d’Inverno (2012) describe a similar approach to computer-assisted composition, in which the artist makes important aesthetic decisions and exercises selective control over computer-generated material. The purpose of this approach is to extend and reflect on one’s own compositional practice, by ‘disrupting habits’ and discovering new creative possibilities.

## 7. TOWARDS A DISTRIBUTED HUMAN–COMPUTER CO-CREATIVITY

An example of distributed human–computer co-creativity in music is the case of *interactive compositions*, that is, compositions involving real-time interaction and mutual adaptation among musicians and software agents. Interactive compositions account for only a small fraction of applications of interactive music systems, the vast majority of which are designed for human–computer improvisation (Gioti 2017). Nevertheless, they provide a fruitful domain for exploring the creative potential of distributed human–computer co-creativity, by challenging traditional notions of authorship and the binary of composition/improvisation.

Interactive compositions differ from interactive improvisation systems in that the design of the software agents involved in them is not just idiom-specific (e.g., designed for jazz improvisation), but composition-specific and therefore far more idiosyncratic. Interactive compositions usually consist of several interaction scenarios, composed in terms of sonic and interaction affordances. Moreover, they are characterised by a prioritisation of interactivity over compositional control, and therefore *process over product*, manifested in the real-time (human and computational) decision-making involved in them. This emphasis on the process is not new: works by John Cage (1960) and Cornelius Cardew (1967), among many others, are examples of similar process-over-product approaches. What is new in the case of interactive compositions is that creative responsibility is shared among not just human, but also non-human actors. The actions of the latter are the result of an extension of human intentionality through technological intentionality (i.e., the intentionality, or ‘directedness’, of the algorithms themselves, determined by what is computationally possible).

Verbeek (2008) distinguishes between three types of intentionality involved in human–technology relations: *technologically mediated intentionality*, *hybrid intentionality* and *composite intentionality*. *Technologically mediated intentionality* occurs when human intentionality is carried out by technological artefacts; for example, when a pair of glasses is used to enhance human vision. *Hybrid intentionality* occurs when the human and the technological merge; for example, when an artificial valve is implanted to replace a patient’s defective heart valve. Finally, *composite intentionality* is defined as the addition of human and technological intentionality, whereby technological intentionality is understood as the way in which a technological artefact is directed at the world. As an example of this type of intentionality, Verbeek (2008) mentions radio telescopes, which produce images of stars by detecting radiation not visible to humans. Verbeek’s examples of all three types of ‘cyborg’ intentionality are limited to physical objects

and do not include algorithms as technological artefacts. Nevertheless, the concept of a composite or, better yet, *extended intentionality* can be useful in the context of distributed human–computer co-creativity.

In distributed human–computer co-creativity, this *extended intentionality* is the result of compositional intentionality carried out by software agents that utilise machine learning and/or generative, rule-based processes. These software agents are capable of creative decisions within a musical space defined by the (human) composer. In this way, creativity is distributed between the human and the computer. To use Boden’s terminology, the human *defines* the conceptual space and the computer *explores* it. This, of course, implies an asymmetry in the relationship between human and machine, mirroring the disparities between human creativity and computational affordances.

## 8. CONCLUSION

Xu, Wang and Bhattacharya (2010) cite Michalos (1970) to distinguish between two forms of creativity: one that is concerned with ‘problem-solving’ and is goal-driven and one that is concerned with ‘problem creation’ and is impulse-driven. They argue that design research on artificially intelligent systems has focused primarily on goal-directed problem solving, ignoring the problem creation phase, which should logically precede it, and that the design process should first address the ‘why’ and then the ‘how’ (Xu et al. 2010).

Similar criticism can be directed towards automatic composition systems. The latter seem to be based on the assumption that composition is ‘problem-solving’ – the problem being one of style imitation – rather than ‘problem creation’ and, as a result, produce artefacts of limited aesthetic value. That is not to say that research on autonomously creative systems is not valuable or that it cannot produce useful knowledge. It is the aesthetic potential of such research that is questionable, not its epistemic value.

From an aesthetic viewpoint, the disparities between human and computational creativity seem to suggest that an ecosystemic approach to creativity, encompassing both humans and machines, is potentially more advantageous, since it allows for a dialogical relationship between human expertise and computational affordances. In distributed human–computer co-creativity, high-level aesthetic decisions are made by humans, while non-human agency is understood as an extension of human intentionality, enabling new types of human–technology interaction and the redefinition of conceptual spaces and artistic practices. The premise behind this approach is that art creation can benefit from a synergy between human and machine intelligence, in which both humans and machines do what they do best.

## Acknowledgements

The author would like to thank Simon Emmerson, Marko Ciciliani and Gerhard Eckel for their critical readings of this manuscript. This research was funded by the Austrian Science Fund (FWF): AR 483-G24.

## REFERENCES

- Ariza, C. 2009. The Interrogator as Critic: The Turing Test and the Evaluation of Generative Music Systems. *Computer Music Journal* 33(2): 48–70.
- Ashby, W. R. 1964. *An Introduction to Cybernetics*. London: Methuen (orig. pub. Chapman and Hall, 1956).
- Bentley, P. J. and Corne, D. W. (eds.) 2002. *Creative Evolutionary Systems*. London: Academic Press.
- Boden, M. 2004. *The Creative Mind: Myths and Mechanisms*. London: Routledge.
- Boden, M. 2007. Creativity: How Does It Work? In M. Krausz, D. Dutton and K. Bardsley (eds.) *The Idea of Creativity*. Boston: Brill, 237–50.
- Boden, M. 2010. *Creativity and Art*. Oxford: Oxford University Press.
- Bown, O. 2012. Generative and Adaptive Creativity: A Unified Approach to Creativity in Nature, Humans and Machines. In J. McCormack and M. d’Inverno (eds.) *Computers and Creativity*. Berlin and Heidelberg: Springer, 361–81.
- Cage, J. 1952. *4'33"*. Leipzig: Edition Peters, score EP6777.
- Cage, J. 1960. *Variations I*. New York: Henmar Press.
- Cardew, C. 1967. *Treatise*. Buffalo, NY: Gallery Upstairs Press.
- Cohn, G. 2018. AI Art at Christie’s Sells for \$432,500. *New York Times*, 25 October 25. [www.nytimes.com/2018/10/25/arts/design/ai-art-sold-christies.html](http://www.nytimes.com/2018/10/25/arts/design/ai-art-sold-christies.html) (accessed 15 May 2019).
- Collins, N. 2012. Automatic Composition of Electroacoustic Art Music Utilizing Machine Listening. *Computer Music Journal* 36(3): 8–23.
- Cope, D. 1992. Computer Modelling of Musical Intelligence in EMI. *Computer Music Journal* 16(2): 69–83.
- De Jager, H. 1972. Some Sociological Remarks on Innovation in Music. *International Review of the Aesthetics and Sociology of Music* 3(2): 252–8.
- Dorin, A. and Korb, K. B. 2012. Creativity Refined: Bypassing the Gatekeepers of Appropriateness and Value. In J. McCormack and M. d’Inverno (eds.) *Computers and Creativity*. Berlin and Heidelberg: Springer, 339–60.
- Engelbart, D. C. 1962. *Augmenting Human Intellect: A Conceptual Framework*. Report AFOSR-3233. Menlo Park, CA: Stanford Research Institute.
- Galanter, P. 2012. Computational Aesthetic Evaluation: Past and Future. In J. McCormack and M. d’Inverno (eds.) *Computers and Creativity*. Berlin and Heidelberg: Springer, 255–94.
- Gioti, A. M. 2017. Machine Listening in Interactive Music Systems: Current State and Future Directions. *Proceedings of the 2017 International Computer Music Conference*. Shanghai: ICMA, 216–20.
- Ito, J. 2016. Extended Intelligence. <https://pubpub.ito.com/pub/extended-intelligence> (accessed 15 May 2019).
- Jones, D., Brown, A. R. and d’Inverno, M. 2012. The Extended Composer. In J. McCormack and M. d’Inverno (eds.) *Computers and Creativity*. Berlin and Heidelberg: Springer, 175–204.
- Latour, B. 2005. *Reassembling the Social: An Introduction to Actor-Network Theory*. Oxford: Oxford University Press.
- Licklider, J. C. R. 1960. Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics* 1: 4–11.
- Manzelli, R., Thakkar, V., Siahkamari, A. and Kulis, B. 2018. An End to End Model for Automatic Music Generation: Combining Deep Raw and Symbolic Audio Networks. *Proceedings of the 6th International Workshop on Musical Metacreation (MUME 2018)*. <http://musicalmetacreation.org/mume2018/proceedings/Manzelli.pdf> (accessed 27 November 2019).
- Martindale, C. 1999. Biological Bases of Creativity. In R. J. Sternberg (ed.) *Handbook of Creativity*. Cambridge: Cambridge University Press, 137–52.
- McCormack, J. 2012. Creative Ecosystems. In J. McCormack and M. d’Inverno (eds.) *Computers and Creativity*. Berlin and Heidelberg: Springer, 39–60.
- Meyer, L. B. 1983. Innovation, Choice, and the History of Music. *Critical Inquiry* 9(3): 517–44.
- Michalos, A. C. 1970. Book Review: *The Sciences of the Artificial*, by Herbert A. Simon. *Technology and Culture* 11(1): 118–20.
- O’Hear, A. 1995. Art and Technology: An Old Tension. *Royal Institute of Philosophy Supplement* 38: 143–58.
- Oliveros, P. 1974. *Sonic Meditations*. Baltimore: Smith Publications.
- Reich, S. 1968. *Pendulum Music*. Vienna: Universal Edition, score UE16155.
- Van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., et al. 2016. WaveNet: A Generative Model for Raw Audio. arXiv preprint arXiv:1609.03499.
- Varèse, E. and Wen-Chung, C. 1966. The Liberation of Sound. *Perspectives of New Music* 5(1): 11–19.
- Verbeek, P.-P. 2008. Cyborg Intentionality: Rethinking the Phenomenology of Human-Technology Relations. *Phenomenology and the Cognitive Sciences* 7(3): 387–95.
- Xu, D., Wang, Y. and Bhattacharya, S. 2010. Thinking Beyond Means-Ends Analysis: The Role of Impulse-driven Human Creativity in the Design of Artificially Intelligent Systems. In D. N. Hart and S. D. Gregor (eds.) *Information Systems Foundations: The Role of Design Science*. Canberra: ANU E Press, 213–32.