INSAM
JOURNAL OF CONTEMPORARY MUSIC, ART AND TECHNOLOGY

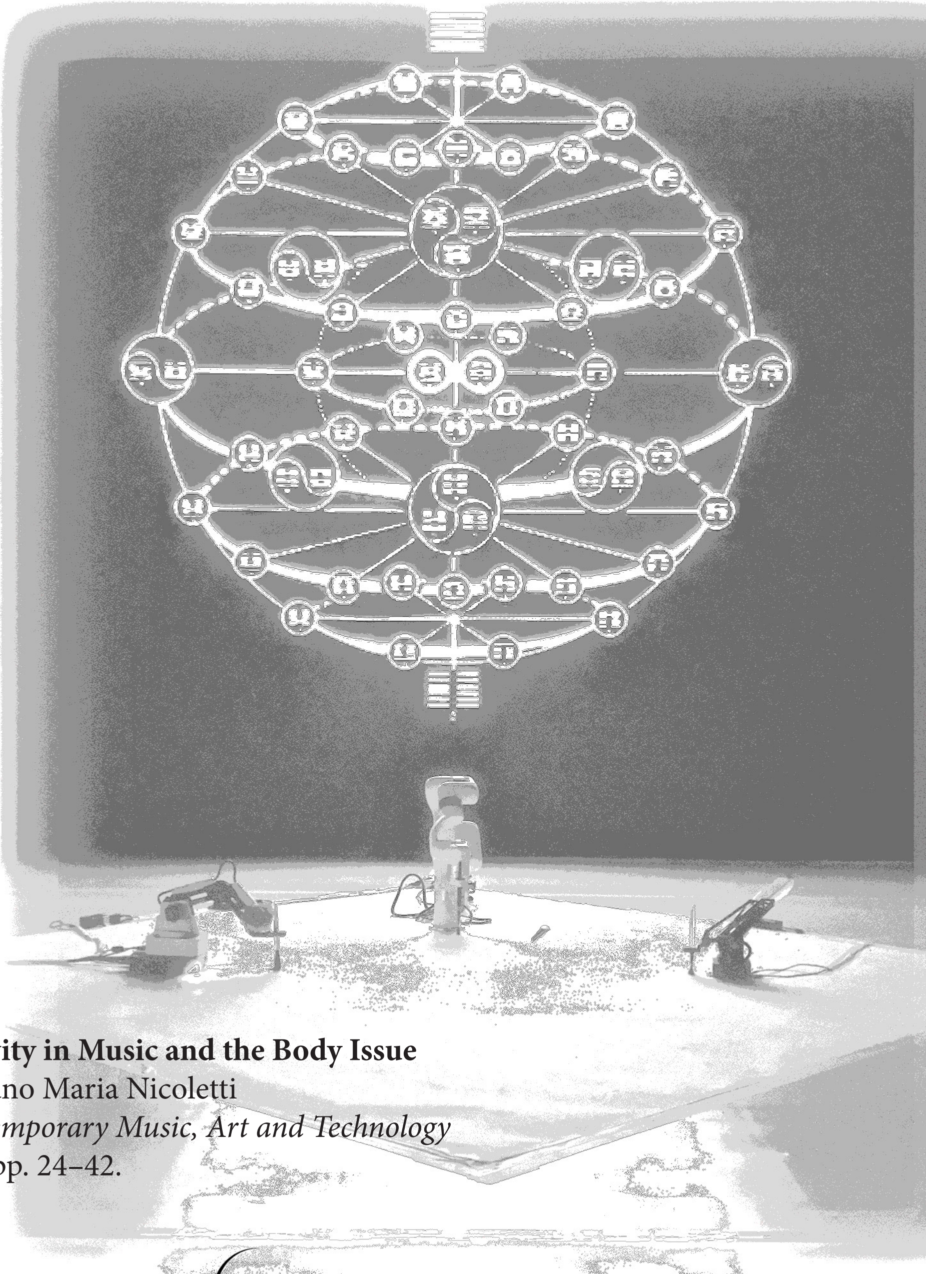

Of Flesh and Steel: Computational Creativity in Music and the Body Issue

Mattia Merlini and Stefano Maria Nicoletti

INSAM Journal of Contemporary Music, Art and Technology
No. 4, Vol. I, July 2020, pp. 24–42.

INSAM

UDC 004.8:78

Mattia Merlini*
University of Milan
Milan, Italy

Stefano Maria Nicoletti**
University of Twente
Enschede, Netherlands

OF FLESH AND STEEL: COMPUTATIONAL CREATIVITY IN MUSIC AND THE BODY ISSUE

Abstract: Could machines ever take our place in the creation of art, and particularly music? The outstanding results of some well-known AIs (e.g. EMI, Flow Machines) might make us believe that this is the case. However, despite this evidence it seems that machines present some intrinsic limits both in creative and non-creative contexts (already highlighted by John Searle and the debate around mechanism). The arguments of this paper are centred around this very belief: we are convinced that the utopian claims regarding all-round machine intelligence are not plausible and that our attention should be directed towards more relevant issues in the field of computational creativity. In particular, we focus our attention on what we call the "body issue", i.e. the role of the body in the experience and creation of music, that we consider problematic for the idea of a truly creative machine (even if we take into consideration weaker renditions of artificial intelligence). Our argument is based on contemporary findings in neuroscience (especially on embodied cognition) and on the theories of Maurice Merleau-Ponty and Roland Barthes.

Keywords: artificial intelligence, computational creativity, mirror neurons, embodied cognition, embodied simulation, body, creativity, performance

1. Introduction

Technological advancement is – as of today – primarily taken for granted: Smart Assistants help us schedule our appointments in our diaries, algorithms suggest what music we should listen to and what we may want to buy next. Our lives are shared with other people at least as much as they are with computational devices and – as

* Author's contact information: mattia.merlini@studenti.unimi.it
** Author's contact information: s.m.nicoletti@utwente.nl

philosopher Luciano Floridi puts it – "we are increasingly delegating or outsourcing to artificial agents our memories, decisions, routine tasks, and other activities in ways that will be progressively integrated with us" (Floridi 2016, 94). The very fact that so many tasks – often considered to be exclusively pertaining to humans – are being delegated to artificial agents challenges our intuitions regarding humanity and its core characteristics. We have observed this shift in perception in the past, for example, with the game of chess: while mastering chess was arguably considered a definitive mark of intelligence before the advent of computers, supporting this consideration is not so easy today, where computational devices can challenge each other to establish supremacy (Silver et al. 2018).

On a par with intelligence, creativity is another trait considered to be exclusive and characteristic of human beings. This exclusivity, concerning creativity, seems to be challenged by the advent of computational devices. In this paper, we reflect upon the relationship between human agents and music, both in the realms of music experience and music creation, in light of historical and contemporary findings in the field of computational creativity. More precisely, we are convinced that there is a hiatus between machine-based (so called) creativity and a human-based one and that this gap could be represented by at least four issues: the "body issue", the "social issue", the "experiential issue" and the "consciousness issue". We have previously sketched these issues (Merlini & Nicoletti 2020), under the hypothesis that these problems are among the most meaningful in the debate concerning the differences between human and machine-based creativity. In doing so we have argued that the concerns regarding all-powerful Artificial Intelligences – supposedly destined to dominate us in every context – are not sufficiently plausible and that we should not grant them priority over the aforementioned issues (cf. ibid., the Chinese Room argument in Searle 1980 and the considerations regarding Strong and Weak AIs, as well as mechanism, presented in Aldini, Fano & Graziani 2016, Beccuti 2018 and Gödel 1951).

In this paper we would like to outline a brief introduction to the most widespread techniques in computational music generation (Section 2), followed by a summary of the aforementioned issues (Section 3). We then shift our attention towards the body issue – that is, the importance of having a body in order to experience and create music. In this section (Section 4), we argue in favour of this centrality by primarily taking into account some contemporary findings in neuroscience, introduced by the philosophical positions of Roland Barthes and Maurice Merleau-Ponty. In Sections 5 and 6 we consider some examples and some possible objections as a conclusion.

2. Computational Efforts and Music Generation

The idea of bringing creativity and computation together is, as a matter of fact, older than the invention of modern computers. One of the very first references to this possibility dates back to Charles Babbage and Ada Lovelace, who were convinced that their Analytical Engine – under certain assumptions – "might compose elaborate and scientific pieces of music of any degree of complexity or extent" (Babbage 1889, 23). The idea of integrating computers and creativity was then taken up by the theorists of the early computer era, such as Alan M. Turing who, around a century after Babbage's ideas:

> was producing (as a joke) programmed love-letters on Manchester's MADM computer; and haikus would soon be generated on Cambridge's EDSAC machine. Even more to the point (or so it might seem), "creativity" was identified as one of the chief goals in the document planning the *Dartmouth Summer School* of 1956. That meeting was where artificial intelligence was officially named (Besold *et al.* 2015, v).

From these early days, the efforts to combine computers and creativity have been manifold and diverse in nature, tackling the fields of visual arts (see Cohen 1995 and Colton 2012), poetry (Colton et al. 2012) and music. The field of music generation – and *Music Generation Systems* (MGS) in particular – has seen the introduction of some of the most notable algorithms developed to combine computation and creativity. The taxonomy described below has been developed by Carnovalini and Rodà (2020, 8-12), a taxonomy which is in turn based on the work of Fernandez and Vico (2013). We briefly present the seven categories in order to capture the most widespread methodologies in the field of computer-generated music, as well as some of the most notable examples from an historical standpoint:

1. *Markov Chains*
2. *Formal Grammars*
3. *Rule/Constraint based systems*
4. *Neural Networks/Deep Learning*
5. *Evolutionary/Genetic Algorithms*
6. *Chaos/Self Similarity*
7. *Agents Based Systems*

The first category addresses *Markov chains*: these are *stochastic processes* – mathematical models that evolve over time in a probabilistic manner – where the outcome of a certain state depends only on the outcome of the previous one (Kemeny & Snell 1976, 1). Suppose we have only three states and the probabilities to transition from one state to the next (from the current generation to the following) are distributed as follows:

		Next generation		
	State	**1**	**2**	**3**
Current generation	**1**	0.65	0.28	007
	2	0.15	0.67	0.18
	3	0.12	0.36	0.52

Table 1

We could then represent this specific process through a *transition diagram* (Kemeny & Snell 1976, 2):

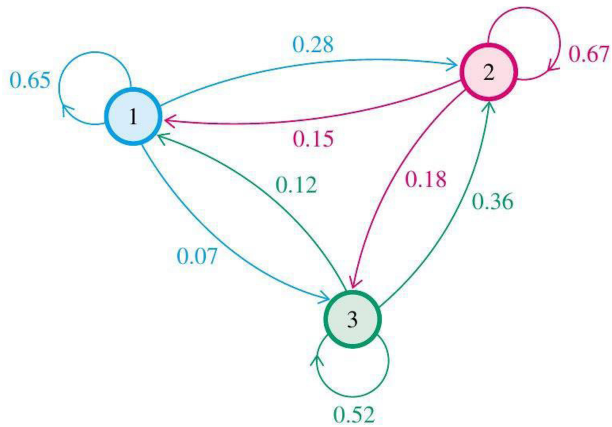

Figure 1. *Transition diagram* (Kemeny & Snell 1976, 2).

Markov chains were chosen by Anderson, Eigenfeldt and Pasquier to propose a generative music system able to compose Electronic Dance Music:

> The Generative Electronic Dance Music Algorithmic System (GEDMAS) is a generative music system that composes full Electronic Dance Music (EDM) compositions. The compositions are based on a corpus of transcribed musical data collected through a process of detailed human transcription. This corpus data is used to analyze genre-specific characteristics associated with EDM styles. GEDMAS uses probabilistic and 1st order Markov chain models to generate song form structures, chord progressions, melodies and rhythms (Anderson et al. 2013, 6).

The second category, *formal grammars*, originates from the work of Noam Chomsky (Chomsky 1957), who introduced the concept of *Generative Grammars:*

> A *Generative Grammar* is composed of two alphabets: terminal symbols and non-terminal symbols (or variables). A set of rewriting rules is given over the union of these two alphabets, that allow to transform variables into other symbols (both variables and terminals). The generated language is the set of all the strings of terminal symbols that can be obtained starting from a special variable chosen as starting point (usually called *S*) and applying any number of rewriting rules in sequence (Carnovalini & Rodà 2020, 9).

One of the most important MGS in computational creativity history is closely related – at its core – to the possibilities granted by formal grammars: David Cope's *EMI*, or *Experiments in Musical Intelligence* (see Cope 1991 and Cope 1992). The combinatorial process that takes place after the analysis of a piece of music is described as:

> The refitting of juxtaposed elements of a work back into logical and musical orders can be enhanced by using *augmented transition networks* (ATNs), a technique developed by researchers in natural-language processing. [...] ATNs can be applied to the recombinant music problem in much the same way as to language: analyze and store musical elements and then reuse them in compositions that vary but have essentially the same musical meaning (variations within a set style) (Cope 1991, 26).

Rule/Constraint based systems, in the third category, differ from generative grammars in being generally unable to produce music out of the blue. They usually rely upon a given input that is subsequently shaped through the application of *rules* and *constraints*: "The inclusion of rules can be implemented in many ways, for example as a final validation step, or to refine intermediate results [...] Constraints can be used to model more abstract features [like tension], rather than explicit music theory rules" (Carnovalini & Rodà 2020, 10). The idea of applying rules and constraints rises pretty early in the history of music generation systems and we can appreciate it in the first two movements of the *Illiac Suite* (cf. Hiller & Isaacson 1958): "The Illiac music was generated from rules defining various styles (including sixteenth-century counterpoint and 12-tone music, and a variety of dynamics and rhythms), sometimes combined with tone-pairs chosen by chance" (Besold et al. 2015, vi).

The complexity of the systems put in place increases within the category of *Neural Networks/Deep Learning*. A notable advantage over the aforementioned methodologies is that the use of deep learning (or generally, machine learning) techniques can create *generality*:

> As opposed to handcrafted models, such as grammar-based or rule-based music generation systems, a machine learning-based generation system can be *agnostic*, as it learns a model from an arbitrary corpus of music. As a result, the same system may be used for various musical genres. Therefore, as more large scale musical datasets are made available, a machine learning-based generation system will be able to automatically learn a musical style from a corpus and to generate new musical content (Briot et al. 2020, 5).

MidiNet, the model proposed in Yang, Chou & Yang (2017), based on a Convolutional Network, is able to "generate melodies either from scratch, by following a chord sequence, or by conditioning on the melody of previous bars (e.g. a priming melody), among other possibilities." (Ibid., 1).

The fifth category is dedicated to *Evolutionary/Genetic algorithms*. As stated by Carnovalini and Rodà (2020, 11), there are three main prerequisites that need to be satisfied to solve a problem via a *genetic algorithm*:

1. *The ability to generate random but suitable solutions to the problem as a starting population*
2. *A way to evaluate the "fitness" of a solution*
3. *The ability to mutate and recombine those solutions*

In doing so, we could operate a continuous selection on the 'fittest' solutions (whose original pool is randomly generated in a suitable manner) to our problem through every iteration of the algorithm itself. One famous example of an algorithm of this kind is *GenJam*, by John Biles. As explained by the author, *GenJam* is:

> a genetic algorithm-based model of a novice jazz musician learning to improvise. GenJam maintains hierarchically related populations of melodic ideas that are mapped to specific notes through scales suggested by the chord progression being played. As GenJam plays its solos over the accompaniment of a standard rhythm section, a human mentor gives real-time feedback, which is used to derive fitness values for the individual measures and phrases. GenJam then applies various genetic operators to the populations to breed improved generations of ideas (Biles 1994, 131).

The sixth category is *Chaos/Self Similarity*. With these methods, musicians and technicians try to generate music that present some degree of self-similarity, either in structures or melodies. In order to accomplish this objective, one possible strategy is to use Cellular Automata: abstract computational systems which are discrete in nature, being composed of a finite set of simple units, or the *cells* (cf. Berto & Tagliabue 2017). Cellular Automata are particularly powerful computational devices that, given appropriate rules, can emulate a Universal Turing Machine and hence compute anything calculable, if we accept Turing's thesis (see Copeland 2020, Turing 1936, and Church 1936). One of the most famous implementations of Cellular Automata to music generation is *CAMUS*, presented by Eduardo Miranda (Miranda 1993). The first prototype, *CAMUS V1.0*, relied on the combined action of two different automata: the first one, based on John Conway's *game of life* (see Berlekamp, Conway and Guy 1982), was responsible for pitch selection while the second was responsible for orchestration. Despite these efforts, however, the results of the systems based on Cellular Automata are not particularly thrilling, according to Carnovalini and Rodà:

> Cellular Automata tend to generate melodies that are not too pleasing, and often need further human intervention. [...] The usual lack of aesthetic value of the results suggests that this is not a good example of CC but rather a way to explore unusual melodies. For these reasons, these systems are arguably less interesting to AI practitioners (Carnovalini & Rodà 2020, 12).

The seventh and last category is dedicated to the *Agents Based Systems*. A software agent is a somewhat autonomous piece of software which is able to perceive and act towards a certain environment with some specific capabilities (i.e., information gathering, learning abilities or "cooperation (with other agents) in order to perform tasks for their owners", Nwana 1996, 213). This kind of software is of particular interest for music generation especially when we see the presence of multiple agents that are able to interact with each other: such a system is then called a *multiagent system* (Vlassis 2007, 1). An example of multiagent system is *Voyager*, presented by George Lewis:

> [...] the Voyager program is conceived as a set of 64 asynchronously operating single-voice MIDI-controlled 'players', all generating music in real time. Several different (and to some, clashing) sonic behavior groupings, or ensembles, may be active simultaneously, moving in and out of metric synchronicity (Lewis 2000, 34).

Voyager essentially creates a virtual orchestra that is able to improvise in a concordant manner and can play together with a human performer, to whom it can react during the execution. Aside from the specific computational effort, (multi)agent-based systems are also important for their attempt towards "humanising" computational means, an aspect that cannot be overlooked when trying to generate music relatable to humans (Carnovalini & Rodà 2020, 12).

3. The Four Issues

Despite these remarkable achievements in the field of computational creativity, we argue that actual creativity is unlikely to ever be reached by any computational or AI-based entity. As exposed elsewhere (Merlini & Nicoletti 2020), we think there are (at least) four main issues preventing AIs from becoming truly creative. In this paper, we focus on what we call the "body issue", but we will also offer a brief overview of the other three issues to create an input for further research. Such issues arise clearly as we leave reductionist positions behind and stop considering music as an isolated (and primarily written) text. From our perspective, it is clear that music must be understood within a frame including its social and cultural context, but also within its performative and corporeal nature. Such aspects of music are not just interesting add-ons for the experience and creation of music; instead, they add to the very essence of the musical experience itself.

Let us try to delve deeper on the first issue – the social aspect of music. The "human" quality of music does not only concern what is commonly called "emotion". Indeed, some AIs can produce music that creates emotional experiences amongst human listeners – and surely, there is no limitation that prevents us from believing that further developments in this field could lead to convincing results. Yet the meaning of music largely depends on social elements that are more difficult to experience when a listener is not a member of a particular society with its own specific culture and history. Computers might be able to create a song "in the style of The Beatles", but that song will always ignore the social value that an actual Beatles' song contains. Music and its meanings always rise from tangible situations. People gather round certain kinds of music and its perceived quality largely depends on elements that transcend its formal features (Spaziante 2007, 33). Could we ever feel the same with computer-generated music? And will the creator act in a way to actually respond to the social needs of the time? One might agree with David Cope (2001, 335) as he argues that music is an independent entity, since the only tangible thing we have is the final product, but this is a highly debatable claim.

Second, we have the experiential issue, which is to some extent consequent to the above. It is easy to argue that computers can replicate the final product without encountering any relevant obstacle – especially when discussing avant-garde hyper-rationally (or randomly) built music, in which the presence of the author attempts to commit its definitive suicide. Yet what we would miss here is not only the idea and concept behind art, but also the very human process that brings us to that. Iannis Xenakis, for instance, remarks (2003) how much his work was inspired by very personal choices and conditions, such as cultural roots, interests, ideals, and the rejection of much of the avant-garde music of his time. Without these human experiences, there is no adequate nest for music to become fully meaningful. The only experience that a computer can gain – as it lacks intentionality and connections with the world – is manipulating 0s and 1s. As Jean-Jacques Nattiez (2007) highlights (focusing on the "aesthesic" aspect of creativity, which is different from the aesthetic element and has to do with the way listeners experience music), the human experience sets the conditions for musical creation: the composer's choices spring from situational factors, discourses and personal experiences, such as their opinion of other composers and their music. Such an interaction with the world is not available to AIs.

This leads us directly to the third issue, which focuses on consciousness. This has nothing to do with emotion, intention or self-awareness, which would lead to an unnecessary amount of speculation. What we are writing about here concerns consciousness in a more phenomenological sense (Brentano 1874), as the human ability to intentionate the world and have a trade with it. Although Searle's position is far from the phenomenological tradition, the basic problem behind strong artificial intelligence – as presented in his Chinese room argument (Searle 1980) – is similar to our point here, as it deals with the incapability of machines of having a "qualitative" experience of the world. From such a perspective, weak AIs seem to be the only possibility, and we should not fear to ever be replaced by truly creative artificial agents. However, there is also an additional reason, on which we will focus in the next section.

4. The Body Issue

Consciousness is tightly connected with the body, so the corporeal side of musical experience cannot be put aside when talking about what makes the human relationship with music so special – which is, after all, what we have been doing up to this point. The strong connection between consciousness and body is something that we can already find in the thoughts by Maurice Merleau-Ponty, who took the phenomenological claims to a whole new level by emphasising the role of our body as living flesh in our embodied experience of the world (Merleau-Ponty 1945). His critiques are directed to abstract conceptions of consciousness that do not take into serious account the carnality of our existence – something that can be found as well in the words of another French author: Roland Barthes. In this case, Barthes (1977) highlights the forgotten bodily side of music, in the form of *musica practica*, that is the bodily sensation of actually playing a piece of music. Everyone can relate to this when thinking of Barthes' idea of the "grain of the voice" (Ibid., 49-55), or the physical side of vocal music, the quality of its sound that conveys all the corporality of the emission, the vibration of the vocal cords and the effort of the larynx. The roughness or the frailty of the vocals in a song, the singer's impetus or grace, are all qualities we can easily sense from music, and which mean a lot to us since we all know, to some extent, how it feels like to sing. We do understand such values as meaningful and can later make creative use of them, only because we own (or we are) a body and can relate to what music-making bodies communicate. This is also valid in the case of instrumental performances. After all, is it not true that popular music (and not necessarily only that) is usually composed by actually playing an instrument (Moore 2001, 56-60), that offers specific affordances (Gibson 1979) and carries a whole set of specific body shapes and physical sensations? In this sense, many creative choices are body-driven and much of the communicative strength of music can spring out of the body knowledge we have of it, using the instrument as a medium for our bodies to imprint their motion in a sonic event.

Of course, what was mainly speculation in the Twentieth Century can now be confronted with scientific findings to add plausibility. Merleau-Ponty and Barthes were lucky in describing such a process so early and many of their intuitions can find fascinating parallels in neuroscientific claims (in Corness 2008, Merleau-Ponty has been related to that context). What we are pointing at is the discovery of mirror neurons and, more specifically, the theorisation of *embodied simulation* (Gallese 2005; Gallese & Sinigaglia 2011; for the relationship with music, see Schiavio et al. 2014). The interest in research on embodied music cognition is growing and there is still much to be understood, but let's try to mention at least some of the major findings that can help us to explain the "body issue" in a more specific way. What we are particularly interested in is the role of embodied simulation when it comes to "understanding music corporally", reconnecting our experience to that of the performer we are listening to – and the meaning that such an experience can have for our comprehension of music and for creativity.

Embodied simulation allows the listeners to feel as if they were actually producing the sound with their own actions, to some extent, by "simulating" via mirror neurons – which activate when passively experiencing a goal-directed action, just as they would when actively producing that action, with no mediation from mental states nor cognitive involvement. Since the action required in performing music "involves the perception of purposeful, intentional and organized sequences of motor acts as the cause of temporally synchronous auditory information" (Overy & Molnar-Szakacs 2006, 236), embodied simulation seems to work with music as well. To further explain this using Overy and Molnar-Szakacs' words: "the expressive dynamics of heard sound gestures can be interpreted in terms of the expressive dynamics of personal vocal and physical gestures" (2009, 492). There is even some evidence of a connection between the experience of music, language and action (sharing the same neural resources), which would then be able to communicate meaning and human affect via embodied simulation (Overy & Molnar-Szakacs 2006), basically putting music (or at least its "motor aspect") at the same level of facial expressions or postures in terms of expressing emotions. New paradigms such as these try to account for the importance of "dealing with music" (Reybrouck 2006, 62) in a concrete way, by focusing on our living bodies and their relationship with musical instruments, which can be conceived as appendices to the body, capable of working as interfaces between us and the world of sound, and as tools for the acquisition of musical body knowledge (Ibid., 66).

Some studies (e.g. Haslinger et al. 2005, Haueisen & Knösche 2001; see Calvo-Merino et al. 2005 for dancing), although not explicitly studying embodied simulation, have demonstrated that a stronger neuronal activity occurs when musicians are listening to music performed with the instrument they are able to play (to the point of stimulating micro-movements of fingers or lips), thus suggesting that a "repertoire of acts" is necessary to fully understand the "physical meaning" of what we are listening to. So, expertise – or at least roughly knowing how it feels to play a certain instrument – seems to be a crucial aspect (Leman 2007, 95-96). Nevertheless, non-musicians can also experience simulation to some extent, not only focusing on the voice, which would seem more obvious as it is the most "human" instrument we have and everyone knows how to sing – this might also explain why vocal music is the most widely appreciated and "understood" by casual listeners. Arnie Cox (2016, 28-29), for instance, explains such phenomenon by recalling the concept of "mimetic subvocalisation" or the rough vocal reproduction of melodic contours not only performed by singers, but also by instrumentalists. Additionally, Cox argues that although one could be in the situation of not knowing how it feels to play an instrument, he/she can always imagine what it would be like to do so (Ibid., 51-52). Although these explanations offer fascinating insights for the resolution of the problem, we argue that more work should be done in this direction. Indeed, since a simulation involving vocalisation is not equal, in terms of sensation, to one involving instrumental playing, the physical feedback might be very different. Moreover, such a conception also prioritises the melodic aspect of music, failing to account for other important aspects. Lastly, the involvement of imagination in the second hypothesis seems to bring into the game those mental states that the very definition of embodied simulation excludes from it. Maybe only a very broad conception of "imagination" (or possibly the "mimetic motion imagery" introduced in ibid., 23 and defined as "not deliberate or conscious"), not involving mental states, can fit this role. The importance of rhythm must also be stressed as a form of musical participation as primordial as vocalisation. If melodic contours can be simulated via mimetic subvocalisation, it is possible to imagine that instrumental parts that are more rhythmically connotated (e.g. drums, rhythm guitar, bass, pizzicato strings etc.) can be roughly simulated by relying on rhythmic abilities (in ibid., 34 we find common manifestations of this: toe-tapping, swaying and dancing to music). Greg Corness acknowledges that there is no actual comprehension of the musical gesture in its physicality, but rather of the intention of the performer – otherwise too many people would not possess sufficient body knowledge to "resonate" with the performer's goal-oriented actions (Corness 2008, 23). Yet we are not convinced by the fact that intention is of primary importance here, as our example will soon explain.

A possible solution to this "expertise issue" is proposed by Overy and Molnar-Szakacs (2009, 493) as they argue that listeners might be able to get to deeper and deeper levels of understanding of musical motion following a precise hierarchy:

1. intention level
2. goal level
3. kinematic level
4. muscle level

Only musicians can truly "resonate" up to the muscle level (with special intensity when listening to music played on the instrument they can play, or at least on instruments from the same family as theirs, see Leman 2007, 97), while a musical novice will not gain access to precise information on any level, but would probably still be able to subvocalise, feel the beat (rhythm and voice) and interpret emotional content according to very basic parameters (e.g. pitch height, speed and intensity).

A computer not only does not have a body (made of flesh and featuring a neural network resembling that of human beings and some animals, thus including mirror neurons), but it also does not own any of the aforementioned (innate?) inclinations towards the production of – say – vocal and rhythmic music. All the communicational strength of such body knowledge is lost, not only because a computer cannot understand music in this way, but also because it is not able to perform music in a way that is meaningful for us. After all, "music is clearly not just a passive, auditory stimulus, it is an engaging, multisensory, social activity" (Overy & Molnar-Szakacs 2009, 489) and, in opposition to a "long tradition of objectivation" within musicology, "music users are biological organisms that have a body equipped with the necessary tools for action, perception and processing at the level of mental operations" (Reybrouck 2006, 60). None of this applies to an AI.

These statements lead us to the final two sections of this paper, in which we will discuss an example and a potential problem, respectively capable of giving a better account of what we are arguing here.

5. An Example: Mono and Tremolo Picking

As music is linked to physicality, it is very difficult to make the reader feel what we are trying to describe in terms of musical experience. Thus, we will briefly analyse a track by the Japanese band Mono – namely 'Cyclone' from The Last Dawn (2014) – to explain the importance of the body issue from a practical point of view, inviting the reader to listen to the song to fully understand what we are talking about. Mono is a post-rock instrumental band best known for their highly melancholic music, described by the guitarist Takaakira Goto as the result of an on-going battle against an overbearing sadness (Chuter 2015, 176-179). This is especially true when it comes to the sibling albums Rays of Darkness and The Last Dawn, both released in 2014 and reflecting witness to one of the toughest moments in Goto's experience (see Chuter 2015). The first represents the dark side (culminating with the highly disturbing 'The Last Rays') while the latter tells the story of a possible redemption.

'Cyclone' is from the second album, yet it is still far from being joyous. We argue that much of the emotional strength and of the meaning of this song owes a debt to the technique Goto uses to play his guitar.

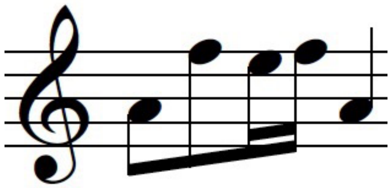

Figure 2. Main melodic contour of the song

The first two minutes of the song present the main chord progression (portrayed by an arpeggiated electric guitar panned to the left) and the main melodic idea that will be carried out for the entire song: a brief and almost circular tune (Figure 2) that perfectly fits with the song's title. From the second minute to the fourth (approximately), the texture becomes more and more dense, as the lead guitar begins to play the same melody in tremolo picking, as typical of much post-rock music (i.e. an alternate picking – up and down – performed non-stop, usually at a very high speed, that confers to the played note a characteristic continuous sound such as experienced in mandolin music). Circularity – which is more likely to be felt as such by a guitarist or an accurate listener, see the "expertise issue" above – takes place at a whole new level, and creates a vivid experience of a cyclone hitting the listener. Yet there is something even more striking happening here: although the average guitarist can perform tremolo picking quite easily, it is still a very "physical" technique and can use high amounts of energy to be performed precisely, especially over a sustained period (as is the case of 'Cyclone'). As Goto plays his part, every guitarist should be able to feel the effort in his/her picking hand, to feel the pain increase alongside the song's dynamics, empathising with the motor act that originally gave birth to the sound he/she is listening to – as the listener is able to decrypt and physically understand its origin thanks to the body knowledge learned from instrumental practice. This is the highest level of simulation: through the sound down to the muscle level (Overy & Molnar-Szakacs 2009, 439), enriching the musical experience with a whole new level of physical and emotive meaning.

This relates negatively with computational creativity if we think of the importance of such an experience when it comes to creating new music, especially in cases in which music is created by directly playing an instrument. The afforded techniques play an important role in musical outcome and, while it is clear that the creator's intention and mood highly influence his/her rendition of the performance, it may also be plausible that composers deliberately or subconsciously choose how to play their music by relying on a "vocabulary of motor acts" which, at least within the same culture, will likely be "felt" in a certain way. Also, staying on a more basic level, it is evident that having a body influences how music is created and how it will be understood by the listener. Although music speaks through sound, it also says a lot "through the body". While computers can (re)produce sounds, they might have a problem with the body.

6. A Problem: Is Electronic Music Cold?

Amongst some music lovers, there may be a widespread belief about electronic music being "cold" when compared to – say – rock or classical music. From a perspective similar to that we have endorsed in this paper, it is plausible that such a claim could find scientific confirmation, given that electronic music often relies on sounds that are not only synthetic, but also triggered by accurately programmed devices involving no human corporality in this process. Since the production of electronic sounds is disembodied (although this is not always the case), there seems to be no space for any kind of embodied simulation. Scholars, who widely acknowledge this issue (e.g. Overy & Molnar Szakacs 2009, 489n; Corness 2008, passim; Reybrouck 2006, 67; Leman 2007, 98; Cox 2016, 37, 212), have tried to suggest different kinds of answers in order to explain why this might not be the case. We challenge the reader to listen to a track such as 'Emerald Rush' by Jon Hopkins without feeling anything connoted in a (very) physical way.

Reybrouck (2006, 67) and Leman (2007, 98) apparently treat electronic-generated sounds as acousmatic sounds (i.e. sounds of which the origin is unknown to the listener). While the first emphasizes the problematic side of the issue, leaving it widely open (at least from the perspective that is of our interest here), Leman later recalls (Ibid., 112) the theories of Theodor Lipps (1903), which, although not particularly linked to the problem of electronic sounds in Leman's argumentation, can suggest a possible solution. Simply put, Lipps argues that we can empathise with the shape of objects, projecting on them what that shape makes us feel like. For instance, a sharp object could make us project on it an upsetting feeling, a certain melodic contour or a specific timbre could recall corporeal articulations or situations. Empathy is still involved here, but in a different way from what happens with embodied cognition. Nevertheless, such a perspective could give us some insights to explain the corporal value of electronic sounds in this direction. Our perplexity here is about the possibility that analysis could easily shift from a context in which the carnal factor is central, to one in which more abstract conceptions of emotivity could become too important – which is something that we do not think will ultimately be able to answer the initial question.

Greg Corness recalls the electronic issue in his 2008 paper, and one of his main questions deals with the problem of gaining different feedback from one person using the computer to write an email in his office and another one using it to perform a set on the stage. Although disembodiment is explicitly addressed (Ibid., 21) as one of the main (if not the main) aspects of the issue, the solution focuses on the context, namely addressing the ability of mirror neurons to "deduct" the agent's intention from the context itself (Ibid., 23). Again, this does not deal with the original problem, and we wonder if that is possible at all. Arnie Cox (2016, 37) describes a hierarchy of sounds from the ones that are the easiest to feel a relation with, to the ones that are less immediate:

1. sounds produced by instruments in which our hands and mouths are directly involved (e.g. voice, hand drums, guitar)
2. sounds produced with the mediation of sticks, keys, bows and such (e.g. drums, piano, violin)
3. electronic sounds created via hand controllers (e.g. keyboard synth)
4. electronic sounds produced and modified via real-time controllers (e.g. knobs and sliders in synthesisers)
5. sounds produced by the playback of recorded music (e.g. musique concrète)
6. incidental human sounds (e.g. Cage's '4'33"')
7. sounds not made by humans (e.g. birdsongs).

The general claim that we can understand from this is that the human touch (and body) can be more or less responsible for the produced sounds. Some of the sounds in which such a responsibility is not evident may sound "enjoyable for some listeners, and disconcerting and unenjoyable for other listeners" (Ibid., 212). That being said, Cox's experience suggests that, while some sounds may "resist" to what he calls "mimetic participation" – and it does not have to be electronic music, since Cox mentions 'Atmosphères' by Ligeti as an example – there is never an occasion in which it comes to zero (Ibid., 48). Electronic sounds may be less "corporal", but somehow still possible to embody. Although this may be proved true by further research and experiments, we cannot avoid asking ourselves if it is true that we get less physical feedback from electronic music.

As is, this issue remains open to this day, and its resolution could lead to interesting developments from the perspective presented here. Understanding how we experience an embodied feedback from electronic sounds may explain why electronic music does not sound as "alien" and "cold" to us, but may also open the way to a reconsideration of the capabilities of AIs. Nothing destructive, because the absence of a body like ours will still prevent computers from physically understanding music and the ability to create something meaningful as a consequence.

The theoretical overview presented in this paper – although not free from problems that have to remain unsolved, for now – argues in favour of a perspective in which the role of the perceiving body is paramount in the process of not only experiencing music, but also creating it, despite the great achievements obtained in the field of computational creativity. This is why we argue that the "body issue" (and the other aforementioned ones) represents an intrinsic limit to computational creativity.

List of References

Aldini, Alessandro, Vincenzo Fano, and Pierluigi Graziani. 2016. "Alcune note sui teoremi di incompletezza di Gödel e la conoscenza delle macchine." *Tra Linguistica e Intelligenza Artificiale*. Università Degli Studi Di Pavia.

Anderson, Christopher, Arne Eigenfeldt, and Philippe Pasquier. 2013. "The generative electronic dance music algorithmic system (GEDMAS)." *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*, Northeastern University.

Babbage, Charles. 1889. *Babbage's Calculating Engines: Being a Collection of Papers Relating to Them, Their History and Construction*. Cambridge: Cambridge University Press.

Barthes, Roland. 1977. *Image, Music, Text*. London: Fontana Press.

Beccuti, Francesco. 2018. "La Disgiunzione Di Gödel." *APhEx* 18.

Berlekamp, Elwyn R., John H. Conway, and Richard K. Guy. 1982. *Winning Ways for Your Mathematical Plays, Vol. 2*. London: Academic Press.

Berto, Francesco and Jacopo Tagliabue, "Cellular Automata", *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/fall2017/entries/cellular-automata/.

Besold, Tarek R., Marco Schorlemmer, and Alan Smaill. 2015. *Computational Creativity Research: Towards Creative Machines*. Paris: Atlantis Press.

Biles, John A. 1994."GenJam: A genetic algorithm for generating jazz solos." In *ICMC* 94: 131–137.

Brentano, Franz. 1874. *Psychologie vom empirischen Standpunkt*. Hamburg: Meiner.

Briot, Jean-Pierre, Gaëtan Hadjeres, and François-David Pachet. 2020. *Deep Learning Techniques for Music Generation*. Cham, Switzerland: Springer.

Calvo-Merino, B., D.E. Glaser, J. Grèzes, R.E. Passingham, and P. Haggard. 2005. "Action Observation and Acquired Motor Skills: An FMRI Study with Expert Dancers." *Cerebral Cortex* 15 (8): 1243–49.

Carnovalini, Filippo, and Antonio Rodà. 2020. "Computational Creativity and Music Generation Systems: An Introduction to the State of the Art." *Frontiers in Artificial Intelligence* 3.

Chomsky, Noam. 1957. *Syntactic Structures*. The Hague: Mouton.

Church, Alonzo. "An unsolvable problem of elementary number theory." *American journal of mathematics* 58, no. 2 (1936): 345–363.

Chuter, Jack. 2015. *Storm Static Sleep: A Pathway through Post-Rock*. London: Function Books.

Cohen, Harold. 1995. "The further exploits of AARON, painter." *Stanford Humanities Review* 4 (2): 141–158.

Colton, Simon, Jacob Goodwin, and Tony Veale. 2012. "Full-FACE Poetry Generation." *ICCC* 2012, University of Dublin.

Colton, Simon. 2012. "The painting fool: Stories from building an automated painter." In *Computers and creativity*, edited by Jon McCormack & Mark d'Inverno, 3–38.

Berlin, Heidelberg: Springer.

Cope, David. 2001. *Virtual Music: Computer Synthesis of Musical Style*. Cambridge, Mass: MIT Press.

Cope, David. 1992. "Computer modeling of musical intelligence in EMI." *Computer Music Journal* 16 (2): 69–83.

Cope, David. 1991. "Recombinant music: using the computer to explore musical style." *Computer* 24 (7): 22-28.

Copeland, B. Jack, "The Church-Turing Thesis", *The Stanford Encyclopedia of Philosophy* (Summer 2020 Edition), Edward N. Zalta (ed.), forthcoming https://plato.stanford.edu/archives/sum2020/entries/church-turing/.

Corness, Greg. 2008. "The Musical Experience through the Lens of Embodiment." *Leonardo Music Journal* 18: 21–24.

Cox, Arnie. 2016. *Music and Embodied Cognition: Listening, Moving, Feeling, and Thinking*. Bloomington, Indianapolis: Indiana University Press.

Fernandez, Jose D., and Francisco Vico. 2013. "AI Methods in Algorithmic Composition: A Comprehensive Survey." *Journal of Artificial Intelligence Research* 48: 513–82.

Floridi, Luciano. 2016. *The 4th Revolution: How the Infosphere Is Reshaping Human Reality*. Oxford: Oxford University Press.

Gallese, Vittorio, and Corrado Sinigaglia. 2011. "What Is so Special about Embodied Simulation?" *Trends in Cognitive Sciences* 15 (11): 512–19.

Gallese, Vittorio. 2005. "Embodied Simulation: From Neurons to Phenomenal Experience." *Phenomenology and the Cognitive Sciences* 4 (1): 23–48.

Gibson, James J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Gödel, Kurt. 1995. "Some Basic Theorems on the Foundations of Mathematics and Their Implications.", 1951. In *Kurt Gödel: Collected Works III*, edited by S. Feferman, J. W. Dawson, C. Parsons, R. M. Solovay and W. Goldfarb, New York: Oxford University Press 3:33–82.

Haslinger, B., P. Erhard, E. Altenmüller, U. Schroeder, H. Boecker, and A. O. Ceballos-Baumann. 2005. "Transmodal Sensorimotor Networks during Action Observation in Professional Pianists." *Journal of Cognitive Neuroscience* 17 (2): 282–93.

Haueisen, Jens, and Thomas R. Knösche. 2001. "Involuntary Motor Activity in Pianists Evoked by Music Perception." *Journal of Cognitive Neuroscience* 13 (6): 786–92.

Hiller Jr, Lejaren A., and Leonard M. Isaacson. 1958. "Musical composition with a high-speed digital computer." *Journal of the Audio Engineering Society* 6 (3): 154–160.

Kemeny, John G., and J. Laurie Snell. 1976. *Markov chains*. New York: Springer-Verlag.

Leman, Marc. 2007. *Embodied Music Cognition and Mediation Technology*. Cambridge, Mass: MIT Press.

Lewis, George E. "Too many notes: Computers, complexity and culture in voyager." *Leonardo Music Journal* (2000): 33–39.

Merlini, Mattia, and Stefano Maria Nicoletti. 2020. "Inhuman, All Too Inhuman: Intrinsic Limits of Computational Creativity in Music." *Riffs: Experimental Writing on*

Popular Music 4 (1): 28-46.

Miranda, Eduardo Reck. "Cellular automata music: An interdisciplinary project." *Journal of New Music Research* 22, no. 1 (1993): 3–21.

Molnar-Szakacs, Istvan, and Katie Overy. 2006. "Music and Mirror Neurons: From Motion to 'e'motion."' *Social Cognitive and Affective Neuroscience* 1 (3): 235–41.

Moore, Allan F. 2001. *Rock, the Primary Text: Developing a Musicology of Rock*. Second Edition. Hants: Ashgate.

Nattiez, Jean-Jacques. 2007. "Alcuni Concetti Fondamentali Di Storiografia Della Musica: Periodizzazione, Spirito Del Tempo, Successione Di Generazioni." *Rivista Di Analisi e Teoria Musicale* 13 (1): 7–35.

Nwana, Hyacinth S. 1996. "Software agents: An overview." *The knowledge engineering review* 11 (3): 205–244.

Reybrouck, Mark. 2006. "Music Cognition and the Bodily Approach: Musical Instruments as Tools for Musical Semantics." *Contemporary Music Review* 25 (1–2): 59–68.

Schiavio, Andrea, Damiano Menin, and Jakub Matyja. 2014. "Music in the Flesh: Embodied Simulation in Musical Understanding." *Psychomusicology: Music, Mind, and Brain* 24 (4): 340–43.

Searle, John R. 1980. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3 (3): 417–24.

Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, et al. "A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go through Self-Play." *Science* 362 (6419): 1140–44.

Spaziante, Lucio. 2007. *Sociosemiotica del pop*. Roma: Carocci.

Turing, Alan Mathison. "On computable numbers, with an application to the Entscheidungsproblem." *J. of Math* 58, no. 345–363: 5.

Vlassis, Nikos. 2007. "A concise introduction to multiagent systems and distributed artificial intelligence." *Synthesis Lectures on Artificial Intelligence and Machine Learning* 1, no. 1: 1–71.

Xenakis, Iannis. 2003. *Musica, architettura*. Milano: Spirali.

Yang, Li-Chia, Szu-Yu Chou, and Yi-Hsuan Yang. 2017. "MidiNet: A convolutional generative adversarial network for symbolic-domain music generation." arXiv preprint arXiv:1703.10847.

OF FLESH AND STEEL: COMPUTATIONAL CREATIVITY IN MUSIC AND THE BODY ISSUE (summary)

Artificial agents are increasingly present in our everyday lives, helping us with our tasks and setting the bar for their “intelligence” higher and higher. Computational creativity is of course one of the main fields in the developments of artificial intelligence, and has been since the very creation of AI. Nowadays, several models for computational creativity are available, and offer very interesting results. For instance, *Markov chains* are used to create Electronic Dance Music, *formal grammars* are featured in AIs capable of creating music recombining chunks from a database, *rule/constraint based systems* are able to generate music applying a given set of rules on pre-existing material, *neural networks* use deep learning to develop music from basic elements, and *genetic algorithms* create suitable solutions for a given context, also being able to mutate and recombine them. Despite such admirable achievements, we argue that the creativity of artificial agents is bound to be limited by a set of (at least) four issues: the "social issue" addresses the importance of the social context in the meaning and value of music; the "experiential issue" focuses on the role that personal experiences and background culture have in the conception of musical ideas; the "consciousness issue" highlights the lack of consciousness as the main limit of creative artificial agents, incapable of having a direct contact with the world. The fourth issue (the "body issue") is linked to the previous one, but focuses on the bodily side of the problem, emphasising the role of the embodied experience of music in the creative process. The discovery of mirror neurons and embodied simulation gives additional persuasion power to the theories of Maurice Merleau-Ponty and Roland Barthes, highlighting the role of the body in the physical and carnal trade with the world. Our body knowledge of music helps us understand music to a deeper extent, without any mediation being required, and the lack of a body made of neurons and flesh seems to constitute a huge limit for computational creativity. Such a comprehension seems to require a “vocabulary of acts”, i.e. an expertise in playing an instrument, to fully take place, so we must better understand to what extent novices can understand music physically, and if there are instruments (e.g. drums, voice) that are more likely to resonate with everyone’s competences. Another important problem concerns electronic sounds that are often created without the implication of direct bodily involvement by the musicians. More work needs to be done to explain why electronic music is apparently no less corporal and vivid as – say – rock or classical music. Since a computer does not own a body, we argue that it is not possible for it to understand music as we do, resonating with the physical motion behind the creation of sounds, but also to use that body knowledge to create fully meaningful music – given that compositional choices often take much of their connotations from the physical action behind their production, as our case study of tremolo picking should demonstrate.

Article received: March 10, 2020
Article accepted: May 20, 2020
Original scientific paper