



**HAL**  
open science

# Artificial Intelligence in Music and Performance: A Subjective Art-Research Inquiry

Baptiste Caramiaux, Marco Donnarumma

► **To cite this version:**

Baptiste Caramiaux, Marco Donnarumma. Artificial Intelligence in Music and Performance: A Subjective Art-Research Inquiry. Handbook of Artificial Intelligence for Music: Foundations, Advanced Approaches, and Developments for Creativity, inPress. hal-03046229

**HAL Id: hal-03046229**

**<https://hal.science/hal-03046229>**

Submitted on 8 Dec 2020

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Artificial Intelligence in Music and Performance: A Subjective Art-Research Inquiry

Baptiste Caramiaux<sup>1</sup> and Marco Donnarumma<sup>2</sup>

<sup>1</sup>Université Paris-Saclay, CNRS, Inria, LRI

<sup>2</sup>Independent Researcher

## Abstract

This article presents a five-year collaboration situated at the intersection of Art practice and Scientific research in Human-Computer Interaction (HCI). At the core of our collaborative work is a hybrid, Art and Science methodology that combines computational learning technology – Machine Learning (ML) and Artificial Intelligence (AI) – with interactive music performance and choreography. This article first exposes our thoughts on combining art, science, movement and sound research. We then describe two of our artistic works *Corpus Nil* and *Humane Methods* – created five years apart from each other – that crystallize our collaborative research process. We present the scientific and artistic motivations, framed through our research interests and cultural environment of the time. We conclude by reflecting on the methodology we developed during the collaboration and on the conceptual shift of computational learning technologies, from ML to AI, and its impact on Music performance.

## 1 Introduction

In many contemporary societies the pervasiveness of technology is constantly expanding. From communication to social networks, digital health and welfare services, every aspect of social life in industrialized societies is being captured by technology with the objective of human enhancement, optimized services, or automated management. Among these technologies, Machine Learning (ML) and the broader field of Artificial Intelligence (AI) received considerable attention in the past decades. In this chapter, we use an autoethnographic approach to present and discuss the hybrid methodology that we developed in five years of collaborative research across computation, science and art. The analysis we offer

here combines insight from academic research in Human-Computer Interaction (HCI), in particular body-based interaction, and from aesthetic research in the performing arts to inspect and question the role of ML and AI in our practices.

The past decade has shown that data-driven learning-based algorithms can succeed in many tasks that were unthinkable not too long ago. Since the deep learning breakthrough in 2012 [44], these algorithms have been shown to recognize images as well as humans do, to acquire motor control capacities from few observations, and to understand and respond to several kinds of human languages. These breakthroughs prompted major investments in the field, thus accelerating technological advance at an exponential pace (the number of attendees at major conferences or the number of published papers at online repository like *arxiv* show the radical increase of interest in the topic). A consequence of the impressive advance of the field is a frantic race to embed ML-based algorithms in most digital services, without offering ways for users to acquire a basic literacy of the technology. The general paradigm behind the development of these technologies is based on centralized ownership by tech companies or institutions. By interacting with these technologies, people produce large amount of usage data, which, in turn, is used to train the underlying algorithms. This type of technology is therefore detached from peoples control and understanding.

Artists have historically been among the first to question technological innovations (see, for instance, the organization Experiments in Art and Technology, E.A.T., founded in 1967). They have often been early adopters and disruptors of new technological tools, and ML is not an exception [13]. On one hand, many artists use ML to enrich the way they work with their preferred media. On the other, some artists use ML, and in particular AI, to shed light onto certain facets of these same tools which can be invisible or taken for granted by public opinion, media, or institutions.

Our own path, which we elaborate next, has taken different directions throughout the past five years of collaboration. Initially, our approach was to use ML as a tool to design improvised and markedly physical musical performance, exploring the relation between computation and corporeality. Eventually, our approach shifted to utilize and reflect on AI as an actor in a performance, an entity whose functioning can be used to question the understanding of computational intelligence in Western society. As a team doing research in overlapping fields, we tried to combine our respective methodologies in a zone of encounter, where *practice-based research* meets *artistic intervention in research*. The former involves the use of artistic practice as a means of research. The latter entails that a creative, artistic process acts upon scientific research objectives. Naturally, such a hybrid methodology has been explored by others before us [28, 26].

In this chapter, we intend to extract the subjective aspects of our collaborative works and to discuss the methodological perspective that our particular mode of collaboration offers to the study of ML and AI, as well as to the study of their impact on HCI and the performing arts. The chapter is structured as follows. First, we provide general thoughts on combining scientific and artistic practices; these observations will help us give context to the methodology of our collaborative work. The next two sections discuss the conception and development of our collaborative artworks, two hybrid performances of computational music and choreography, *Corpus Nil* and *Humane Methods*. We present the scientific and artistic drives behind each project, highlighting scientific, artistic and cultural contexts. Finally, we provide a closing discussion where we bring together the core findings emerged from our research.

## 2 Combining Art, Science and Sound Research

A dialogue between scientific research and artistic practice is inclined to generate multiple and often contrasting perspectives, rather than producing an agreement. Important insight can emerge from those contrasts, insight that would have been obscured by a unified and monodisciplinary approach. This is the methodological grounding of the present chapter. We do not strive for objective knowledge. Art, even in its practice-based forms, does not offer any provable truth but creates alternate visions and, sometimes, it is called upon to contemplate the dark without drawing a resolutely positive lesson, as Fuller and Goriunova put it [33].

### 2.1 Practice-based Research and Objective Knowledge

The kind of art and science coupling we discuss here is not the use of scientific findings for the creation of an artwork, for that strategy requires an attentive process of translation, as stressed by Stengers [53], and based on our subjective experience it can easily feel constraining. As we will elaborate later when discussing our first collaborative project *Corpus Nil*, valid insight or correct measurements achieved in a laboratory setting may not be useful when applied to a musical performance. The motivation lies in the experimental limitations of particular scientific methods found in HCI and other research fields combining natural science, design and engineering. Contrary to the social sciences, for instance, HCI experiments are often conducted in a laboratory context. In controlled studies, every aspect of such studies is thoroughly directed by the experimenters. Whereas the protocols of such HCI experiments aim to ensure scientific rigour and reproducibility, those same protocols

create the specific context wherein the collected data and related findings are meaningful. Outside of the specified context, data and findings may not be as meaningful, especially when applying them to an unpredictable and ever-changing situation such as a live artistic performance before a public audience.

For example, during a live performance, sensor measurements that have been validated through a particular experiment may very well vary according to a range of factors, including room temperature, magnetic interferences and the like. Those variables - contrary to a laboratory experiment situation - cannot be controlled. More importantly, the performer experiences her own body, a technological musical instrument and the context in subjective ways; during an artistic performance a performer is driven by instincts, desires and feelings that arise from the relation with the audience, other performers and a particular musical instrument. Artistic expression in musical performance emerges from the interaction of human actions and desires with the capabilities of a musical instrument.

This emphasises the need for, on one hand, a careful selection and interpretation of the methods of art and science collaboration, and, on the other, a precise definition of the shared knowledge needed by artist and scientist to operate collaboratively; an entanglement that Roger Malina aptly calls deep art-science coupling [46]. Practice-based research represents a good framework for the exploration of art and science coupling, for it generates knowledge from action, self-reflection and empirical experimentation [11, 6], methods that are equally fitting in HCI as in music and the performing arts. Crucially, we believe that the goal of art and science collaboration is not the production or re-staging of the limits of a normative type of science, but rather is about “working outside current paradigms, taking conceptual risks”, following Malina, so as to mindfully merge disciplines towards new practices of experimentation. This can result in contrasts and disimilarities which must be considered and which, in fact, produce the richness of art and science practice.

## **2.2 Artistic Intervention in Scientific Research**

As technology has become more ubiquitous, spreading from workplaces to everyday life, the field of HCI has adopted new ways to investigate how we interact with such pervasive technology [39, 23]. The so-called third-wave of HCI looks at notions such as experience, emotions, aesthetics and embodiment [9]. The field thus embraces methodological resources from a broader range of fields such as cognitive science, social science, science studies, or the Arts. In this context, artistic intervention encourages research practitioners to look at their subject of study under an alternative hypothesis, through different facets. It is important to underline that, differently from practice-based research, artistic intervention in research

emphasizes the idea of intervention: the artistic process *acts* upon the research objectives. For example, artistic intervention in scientific research, and more specifically in HCI, can help address a problem (or study a phenomenon) in radically new ways [5]. Such as in the case of Fdili Alaoui [28] who, working between HCI and dance, questions the academic culture of HCI through her expertise in interactive dance performance.

In our research, we use artistic intervention to build and study interactions between people and intelligent systems. This methodological choice is motivated by two factors. On one hand, scientific research practice can feel normative and standardised. As researchers, we employ established methods to structure our investigations on a particular question, as well as to disseminate the outcomes of the research. While this standardisation has its pragmatic rationale (for instance, to help its dissemination), it can be detrimental to scientific research practice in many ways, such as limiting exploration or surprise. A detailed articulation of the pros and cons of it goes beyond the scope of this paper. Here we stress how artistic intervention in research can complement standard scientific methods in leaving more room for unexpected challenges and nuanced conceptual questioning. As an example, which will be developed in the following section, our study of gestural expressivity in musical performances has shifted the emphasis from characterizing gestural variations in kinematic space to characterizing variations in muscle coordination space.

On the other hand, the specific field of HCI has its own culture and shortcomings. While the field examines and generates the systems underpinning the technological world with which humans interact daily, its political dimension – that is, the varying balances of power between who produces the technology, who designs it, who uses it, and what socio-cultural impact it may have – is rarely addressed [40, 4, 42]. By avoiding to explicitly address the politics of technological innovation, one risks to contribute to the creation of forms of technological solutionism. Technological solutionism refers to the idea that technology can solve any problem, including issues which may not exist or which ignore both socio-cultural and political contexts [8]. We found that one method to prevent (at least partially) our research from entering a solutionist discourse is to invoke artistic intervention within the process of scientific research. Critical artistic intervention can explicitly provoke self-reflection about the involvement of a particular technology. Through an analytical and self-reflexive engagement, a critical artistic intervention can question the role of technology and its capacities of meaning-making, forcing artists and researchers to face the inner beliefs that ultimately motivated their choices.

In the following, we discuss how we applied the methods and insight described above to two separate research projects, which led to the creation of two artistic performances

and scientific publications. This also serves to illustrate how our collaboration developed iteratively through the intertwining of practice-based research in the art field and artistic intervention in the scientific research.

### 3 Machine Learning as a Tool for Musical Performance

Our first project stemmed from a collaboration started in 2014. The outcomes of this particular project were a performance entitled *Corpus Nil* (premiered in 2016 at ZKM, Center for Art and Media, Germany and still touring internationally) and academic publications [12, 22].

#### 3.1 *Corpus Nil*

The piece is a twenty-minute dance and music performance exploring hybrid forms of identity and musicianship. It is an intense and ritualistic interaction between an autonomous musical instrument, a human body, and sound. Figure 1 depicts a picture from a live performance. The theater space is completely dark. The player (Donnarumma in this case), whose body is partly naked and partly painted in black, performs a tense choreography which gradually morphs his body. Two types of wearable biosensors transmit data from his body to our machine learning-based software. Chip microphones capture sounds from muscles and internal organs (mechanomyogram or MMG) and electrodes capture muscle voltages (electromyogram or EMG). The software uses special filters to generate a description of the amplitude and frequencies of all sounds produced within the performer’s body (between 1-40 Hz), as well as their variations over time. Then, it re-synthesises those sounds by orchestrating a feedback network of twenty digital oscillators. Because the choreography demands slow, subtle and iterative physical movements, the resulting music is equally slow and recursive, mutating across microtonal variations of a minimal set of pitches.

The instrument here intended as a combination of our chosen hardware and software analyses the collected data to learn the nuances of the performer’s movement (muscular tension, gesture abruptness, rate of relaxation). Upon observing a certain nuance, it chooses whether to mute or activate particular oscillators, how to regulate volumes, phases, glissandos and multi-channel diffusion, and how to adjust feedback amounts within the network. The player cannot control the instrument, but can only learn how to affect it and be affected by it. The piece thus discards conventional performer-instrument relationships founded on the performer’s full control of the instrument in favour of an unstable corporeal engagement between the two. Through the rhythm of sound, vibration and light, the performer’s



Figure 1: Picture from the performance *Corpus Nil*. Photo courtesy of ONUK Fotografie. The piece evolves through different levels of muscular tension, torsion and articulations of the upper limbs. Aspects of the movements captured by means of two different muscle sensors placed on the arms.

body and the instrument mutate, physically and conceptually, into something other; an unfamiliar creature defying common definition of the human.

The software used in *Corpus Nil*, collaboratively created by the authors, consists of a set of algorithms linked to one another forming a recursive network. The algorithms have different tasks, such as biosignal analysis, learning of movement nuances (through linear regression and statistical methods), and sound resynthesis. A group of algorithms extracts a set of high-level features from the biosignals collected from the player's body (change rate, damping and spectral centroid). This feature set does not represent the movement per se, for it does not account for the shape of movement in space. Rather, it indicates specific traits, such as muscular force, abruptness of a contraction, and damping of the muscular tissues, that characterise the articulation of the movement. The combination of analysis, learning and sound synthesis algorithms into a recursive network makes the instrument



dynamic enough to provide highly unexpected responses to the player’s movements. This, in turn, forces the performer to adapt the choreography on the fly, establishing thus a continuous and adaptive dialogue with the instrument.

### 3.2 Scientific and artistic drives

Here we present the scientific and artistic contexts, specific to each author, in which our collaboration on *Corpus Nil* took place.

**Gesture Expressivity and Interactive Machine Learning** When we started the collaboration in 2014, my (first author) research focused on capturing movement expressivity to be creatively used in music performance. A hypothesis was that aspects of the physical expressivity of a performer could be observed by analysing intentional variations in the execution of a movement. A system which would capture such variations could then be used to drive sound synthesis engines. However capturing those variations proved not trivial. Our approach at the time, developed with colleagues at Ircam (Paris) and Goldsmiths College (London), relied on a machine learning-based algorithm called Gesture Variation Follower (GVF) [14] which tracks, in real-time, continuous movement variations across space and time. In order to allow artists with radically different movement vocabularies and expressivity to use the system, this was designed to be rapidly calibrated.

This research belongs to a broader line of endeavour where machine learning algorithms are considered as tools to design expressive musical instruments and interactions [30, 15, 36, 29, 34, 32]. In this type of research, the significant advantage of machine learning is that it leverages advances in computation and data analysis to allow musicians to tackle fairly complex musical scenarios [29]. More precisely, instead of writing rules that govern the criteria of interaction between performed movements and resulting sound, one can provide the algorithmic system with demonstrations of those criteria, which will be automatically learned by the system [31]. This approach involves tight interactions between the performer and the ML algorithm, formalised under the discipline of Interactive Machine Learning (IML) [55, 27, 1]. Using machine learning as a design tool seemed a natural approach when considering movement inputs to an interactive system. Writing rules that govern the analysis of movement inputs is clearly unfeasible when we consider that the system is expected to handle a wide scope of complex movements, mostly subjective to each performer. In addition, movement literacy is tacit, which means it cannot be easily formalised through words and lines of code of a programming language in a design context [38]. A machine learning-based approach where one can configure a system by demonstration seems a much

more efficient approach [35]. In addition, such a system can be more inclusive, meaning that it can be used by novices and people with diverse physical abilities, who may not have previous experience in movement analysis or computer science [41].

**Deconstructing a Performer’s Body** By the time our collaboration began, I (second author) had been working for four years in the field of gestural music as a researcher, composer and performer; in particular, I had been developing a musical performance practice known as *biophysical music* [21]. With this term, I refer to live music pieces based on a combination of physiological technology and markedly physical, gestural performance. The approach of biophysical music differs significantly from previous strands of biosignal-based musical performance [43, 49, 54] which are largely rooted in the use of physiological signal as a means of control over electronic or digital interfaces. The practice of biophysical music focuses, instead, on designing ways and tools through which the physical and physiological properties of a performer’s body are interlaced with the material and computational qualities of the electronic instruments, with varying degrees of mutual influence. Through this method, which I term *human-machine configuration* [20], musical expression arises from an intimate and, often, not fully predictable negotiation of human bodies, instruments, and programmatic musical ideas (an anthology of biophysical music projects can be viewed in [19]).

After several years exploring potentialities and limits of gestural music research, I started questioning in both my artistic and scientific processes the kind of human body at the center of conventional gestural music performances. My concern arose from both my own work experience in the field and my participation to the HCI community through conferences and research exchanges, as well as from my study of feminist theory, in particular body theory [7] and disability studies [51]. Soon, it became clear to me that the kind of body shown in the performances of gestural music, and new musical instruments in general, was most times a heavily normalised body, that is, a body fitting particular criteria of normality. This triggered me to investigate how elements of gestural music performance and HCI could be exploited in combination with research on movement and aesthetic to create a different kind of performance practice; a practice that would account for the human body as an ever-changing, fluid entity with multiple potentials and varying forms of embodiment, instead of a static, controlled and controlling subject as established by our societal regime. Thus, drawing on the artistic work of performers such as Kō Murobushi [47] and Maria Donata D’Urso [25], among others, I developed a particular methodology of movement research that emphasised the combined use of unconventional choreographic methods and intense

somatic experimentation towards the physical manipulation and symbolic deconstruction of a performers body.

The movement experiments I conducted at the time were based on a gesture vocabulary composed of complex torsions, flexions, and contractions of shoulders, upper arms, and neck. No lower arm or hands gesture figured in the movement vocabulary I had designed for the piece (which then became *Corpus Nil*); hands, lower arms and even the head were, in fact, hidden from view and rendered purposely useless. Because the literature and tools in regard to this particular mode of gestural interaction are scarce, we decided to explore muscle activity at a very low level – through choreographic methods and resources from biomedical engineering. This, in turn, led us to investigate complementary sensor modality, with the goal to understand physical expression outside of a frame focused on “control”.

### 3.3 Development and Observations

The initial seed of the collaboration was an intuition to investigate applications of interactive machine learning to the analysis of movement in a body-based musical performance. The performance in question, however, relied on the expression of bodily and motion qualities that were uncharted by previous work in the field. This spawned a research of muscle activity, the sensorimotor system and how, through different sensing methods, aspects of expressivity could be observed in human-computer interaction.

**Understanding Movement Expressivity** Initially, the analysis of movement expressivity as conveyed by the specific kind of muscle torsions, flexions and contractions described in the previous section seemed a good challenge for the Gesture Variation Follower algorithm (GVF) [14]. GVF was originally designed to track variations of movement trajectories in space and time, so we explored fundamental questions about movement expressivity at the level of the muscle activation. What is the trajectory of muscle activity in this context? To what extent information on muscle temporal activity is meaningful to a performer, and how muscle temporal activity can be controlled?

We began by experimenting with mechanomyogram (MMG) sensors, capturing the mechanical activity of the muscle, built by Donnarumma in previous work [18]. We organised sessions of data capture and post-analysis. It became quickly clear that MMG data captured dynamic transitions between separate muscle contractions and the resulting trajectories calculated by GVF did not make sense. We added a second sensor based on electromyogram (EMG), which measures the electrical discharges activating the muscle. Here a new problem emerged, following Donnarummas feedback about the resulting movement-sound interac-

tion: controlling EMG trajectories was not aesthetically convincing, for it could not allow a highly dynamic mode of interaction, one of our main musical goal. Therefore, we worked on a new way to analyse muscle biosignals. We developed a different tracking system where the tracked parameters were not linked to the movement trajectory but the to the parameters of an underlying second order dynamical system with damping (symbolically representing muscle dynamic as an oscillating string).

How to define gesture expressivity when the gestures in questions operate through a set of symbolic and aesthetic signs? What kind of expressivity emerges from a performer’s body that is physically constrained by a given choreography? And how can a ML system actively and subtly support a semi-improvisational physical performance of music where the very notion of “control” is put into question? We felt that these questions were to be tackled through a strategy that combined scientific and artistic methods. Our research forked in two parallel streams of investigation. On one hand, we deepened our understanding of expressivity to take into account involuntary aspects of whole-body gestures. Using real-time analysis of complementary muscle biosignals, EMG and MMG, we began isolating a set of features that could describe relative levels of muscular effort without the need to define in advance the type or the timing of a gesture. This would support the open-ended and semi-improvised nature of the choreography that was being developed. On the other hand, we started exploring the design of a computational music system that would not be tied to conventional mapping techniques, but instead would reflect, through sound, levels of effort intensity, movement abruptness and degree of complexity of particular gestures i.e., the amount of muscle groups activated by a given gesture.

We finally conducted a controlled experiment to study gestural expressivity from the point of view of motor control (i.e., the ability of users to control gesture variations), using both EMG and MMG interfaces [12]. The main study brought together 12 participants for a session of 45 minutes each. In these sessions, participants were asked to perform a certain number of gestures (with or without tangible feedback) and then to vary one or more dimensions of expressivity. We showed that the participants consistently characterized dimensions of implicit expressivity (force, tension, etc.) and that the physiological interfaces used made it possible to describe these dimensions. Finally, we showed that participants could control these dimensions under certain conditions linked to constraints from the laws of motion.

**Highlighting algorithm’s limitation with IML** GVF was designed to be quickly trained (or calibrated) by providing one example for each movement that the system has

to recognise, and its capacity to track movement variation was meant to give to users a means to explore a movement space. However its use in a tight interaction loop with the performers body during experiments and rehearsals felt like a failure. One major factor was that the system felt restrictive rather than extensive. The movement possibilities seemed to shrink rather than extend. In the typical IML workflow, the performer is engaged in a close interaction loop with the movement-based machine learning system. However, if the algorithm has low-capacity (failing to handle a wide range of input movements) with respect to the complexity of the input movement, the performer may ultimately adapt to the system limitations, hence constraining her own movement qualities. Consequently, this adaptation would make the interaction converge towards simplistic movements and interactions.

This phenomenon may have also been emphasized by the absence of clear explanations about the system’s prediction behavior. Incorrect analysis of the GVF were hard to grasp by the performer (second author). Typically, the algorithm could have failed because the beginning of the movement is too different from the pre-recorded ones, but this information was not clearly fed back to the performer. As most temporal gesture recognition systems, GVF recognises a gesture (or movement) from a starting point until an end point. Establishing what defines the beginning of a movement and its ending point during the artistic performance was not trivial and it could definitely not be handled by the algorithm. Moreover, the system was designed to handle continuous and rather slow variations, whereas Donnarummas muscle activity during the choreography presented types of dynamics that could not be tracked by the system. We needed a more flexible approach and thus we developed a new tracking system. This was designed to estimate in real-time the parameters of a physical model that mimicked the performer’s muscular activity, as mentioned above. In doing so, we were not tracking gesture trajectories but movement “regimes”. Thus, this method allowed us to explore alternative representations of movement using the input data at hand.

## 4 Artificial Intelligence as Actor in Performance

In 2019, we initiated a second collaboration, which actively involves a group of 13 people, including artists working with new media, biological materials, and performance, as well as designers, AI and neurorobotics scientists. The project outcome is a dance-theater production entitled *Humane Methods*. The first version of the performance premiered at Romaeuropa Festival and, at the time of writing, the project is currently ongoing. As for the previous collaborative project, we begin by describing the piece and then situating it

within distinct scientific and artistic contexts.

#### 4.1 *Humane Methods*

*Humane Methods* is a dance-theater production exploring the multilayered nature of today's violence. The project departs from the assumption that through the physical and psychological brutalization of people and non-human beings the combination of digital technologies and capitalistic urge has driven the natural ecosystem towards impending destruction. The project then aims to dissect the violence of algorithmic societies, where power structures, knowledge creation and normative criteria become means of manipulation. Being an evening-length production, *Humane Methods* is composed of a multitude of elements, including robotic prosthetic limbs driven by their own neural networks, dead plant specimen and fungi growing on costumes, uncompleted or destroyed architectural structures, chimeric creatures, experiments in trauma, actions as rituals. For the sake of clarity, here we focus on the technological aspects of the piece and how they relate to the techniques and implications of AI. Figure 2 shows a picture from the performance.

In a landscape of ruined nature and technological relics, a community of nine semi-human figures lives a primal and stern existence. They perform a ritual of prayer in a loop, but with every repetition of the prayer, something in their interpersonal relations or in the world surrounding changes. As variations accumulate, the praying loop is drastically altered and the storyline branches out into a constellation of stories, of hints, of untold events, playing out various degrees of intolerance, isolation and exploitation. Throughout the action, a deep reinforcement learning algorithm tries to learn a meaningless sequence of ten digits.

Our motivation in choosing to work with an AI algorithm performing meaningless calculations lies in a conceptual shift from conventional uses of deep learning techniques in music and performance. What interests us is not the capacity of the algorithm to reach its target, but rather the ways in which the inner (obsessive) logic of this type of computation can be made perceivable at an aesthetic and sensorial level to both performers and audience. To that end, we designed an audiovisual system that links the actions of the AI algorithm to synaesthetic patterns of music and light, so that the algorithm speaks, as it were, of its actions and choices through sound and light. The result is a dogged, hypnotic stream of sound and light that inundate the venue: the calculations of the AI are turned into perceivable material, auditory and visual pulsating patterns that literally mark the choreographic actions on stage, as well as the spectators experience of the piece.



Figure 2: Photo courtesy of Giada Spera. Picture from a live performance of Humane Methods at Romaeuropa Festival. Portrayed is one of the iterations of the ritual at the core of the performance. It is possible to observe in detail the lighting generated by the AI.

## 4.2 Scientific and Artistic Drives

**Human - AI Interaction** ML and AI algorithms have known an exponential development in the past decades. Several important works have shown the capacity of learning-based algorithms to outperform humans in specific cognitive tasks such as playing the game Go [52], recognising images [44] or understanding natural language [37].

Simultaneously, the fields of HCI, Humanities, and Social Sciences have started to take seriously the study of the socio-cultural impact of AI and to propose design guidelines or good practices. Several works have recently been published along this line. Among them, some aim to propose a set of guidelines for designers of AI-powered interactive systems [2]. In this context, AI is seen as a tool to improve users services and generally empower humans [10]. As an extension, AI can be seen as a material that can be appropriated by designers to improve user experience in AI-based applications [24]. However, most of these

attempts have highlighted the inherent challenges of this approach due to the difficulty of grasping what AI can and cannot do [56]. In my work (first author), I am exploring how the computational learning mechanisms themselves can become interactive in order to foster exploration and human learning, as we recently explored in the specific domain of sound design [50].

On the other hand, recent works have investigated the complex consequences of the inherent biases in AI-powered applications and their impact when used in health care, education or culture [3, 45, 17]. This last branch of research is a necessary undertaking that I am currently pushing in my recent research interest, which looks at the socio-cultural impact of technology. Therefore, at a scientific level, the collaboration presented in this section was motivated by an intention to investigate how people understands AI technology and which are the founding beliefs supporting particular forms of such understanding.

**Motivations and Artistic Needs** The project *Humane Methods* is a collaborative endeavour that involves, aside from the collaboration between myself (second author) and Caramiaux, a shared authorship between Margherita Pevere [48], myself and video artist Andrea Familiari. In particular, the artistic concept for the piece was jointly created by Pevere and myself, and then furthered through a close collaboration with Caramiaux. As artists working respectively with biotechnology and AI technology, the first question that Pevere and I asked ourselves was how to situate our artistic practices within the current historical period, one characterised by a generally intolerant and polarised socio-political forces, converging with a serious, technologically-driven disruption of the natural environment. This question led our efforts towards an analysis of the multifaceted nature of human violence, as it is exerted between humans as well as it is enforced by humans on the lives of non-human entities and the natural environment.

Turning to the main topic at hand – AI technology and its use in collaborative art and science projects – we chose to deploy AI algorithms in an unconventional way. Instead of using AI to aid human performers in the creation of sound as we had done for *Corpus Nil* or using an AI generative capacity to create a musical automata, we chose to make tangible the “brute force” computing mechanism of AI algorithms. In other words, we chose to highlight how many of the most common learning algorithms being presently developed attempt to reach their targets using a mechanism that is obsessive and raw.



### 4.3 Development and Observations

The collaboration developed through personal exchanges on the general theme of technology and society and particularly on how technology induces violence or empathy on people, to other people and to their environment. We wanted to create an algorithmic behavior that obsessively learned something and which could autonomously explore different ways to reach its goal. The first trials, in terms of scenography, music, light and algorithms happened in June 2019.

**Machine Behaviour** The algorithm designed and implemented for *Humane Methods* is based on deep reinforcement learning, which means that the algorithm explores a high-dimensional continuous space and receives positive rewards when approaching the goal. These rewards are used to exploit certain moves within this space that would help the algorithm reach the goal. In *Humane Methods*, the algorithm's target is an arbitrary set of parameter values which – purposely – holds no meaning. The reward is given by its distance to the target. Its actions are moves in the parametric space. As the algorithm moves in the parametric space, it outputs a list of 10 values indicating the degree of proximity of its calculations in respect to its target on an integer scale between 1 and 3. Each of the 10 values is then mapped to a musical pattern which are predefined arpeggios of acoustic piano samples and a light pulsation pattern. Sonic and light patterns follow relentlessly one after the other, creating the perception of an audiovisual entity morphing continually through a multitude of syncopated variations. This dynamic morphing of sound and light is made more complex by a parallel method linking the algorithm degree of proximity to the target to, on one hand, the loudness of each single musical pattern, and on the other, to the brightness intensity of physical analog lights. The closer a value is to the corresponding target, the quieter is the volume of the related musical pattern and the dimmer is the brightness of the corresponding light pattern.

The algorithm moves within the space are entirely inferred by the learning process. This process consequently makes apparent to the spectators the behaviour of the algorithm. The algorithm becomes thus an actor, a performer on its own. Because neither the human performers nor the technical team behind the piece can control the computation of the algorithm, the latter is the sole entity responsible for the audiovisual dramaturgy of the piece. Interestingly, this role taken by the algorithm made us personify it: we began referring to it by a name, or discussing its contribution to the piece as an actor on its own. Journalists have also asked questions along this line, after seeing the show. This has been all the more striking for us as we were very careful not to anthropomorphize AI, a process

that we find counter-creative for it typically impedes a precise understanding of technology by imposing a human worldview onto something which is not human.

**Controlling Artificial Intelligence and Being Controlled by It** Following our first experiments with the algorithm and, thus, observing the modalities of its interaction with music, light and choreography, it became more interesting to us *how* the algorithm learned rather than what it learned. We tested several strategies to map, more or less directly, the behaviour of the algorithm (the *how*) to the choreography executed by the human performers discarding, in the process, any information regarding what the algorithm learned. In doing so, we soon faced a different challenge: how to find a balance between the control of the AI over the piece's narrative and the influence of the performers over the AI's learning process. We explored different strategies. One was to have the algorithm dictate the start and end points of each choreographic action. This, however, felt narratively constraining and aesthetically poor. Another attempt required the AI to describe the movement qualities with which each loop was to be performed by the human actors. This proved cumbersome, for a reliable, real-time communication between the AI and the performers was too complex to design.

At this point we changed our approach. Instead of mapping the AI's behaviour to the actions on stage, we thought of implementing the opposite strategy, that is, to map the actions on stage to the behaviour of the AI. Thus the choreography was re-conceived to mimic the temporal structure of the AI's learning process. One single action (a hybrid form of gestural prayer) is repeated in a loop, but with each repetition comes a more or less drastic variation. This loop structure mirrors the episodic learning of the algorithm: the algorithm learns through episodes – each can take a few minutes – and each episode adds variations to the next iteration of learning. This solution meant to limit the explicit interaction between actors performance and AI learning process. The interaction between the two elements became thus implicitly intertwined: at conceptual and aesthetic levels the performers' action and the algorithm's learning process were connected, but this connection was not defined by explicit mapping of control signals. While we successfully performed the piece before a public audience in this format, we consider the issue still open. In our view, the piece would have a stronger impact if the performers' actions would be linked *both* implicitly and explicitly to the learning process of the AI.

**Audience Perception of Artificial Intelligence** In June 2019, we had the opportunity to collect subjective feedback from the audience during an open rehearsal at the Centre

des Arts Enghien les Bains (CDA), a co-producer of *Humane Methods* together with RomaEuropa Festival, Italy. Our encounter with the audience in France was important, for it allowed us to open our aesthetics and research to a general public, most of whom had little or no familiarity with the ongoing project or our previous works. Here, we focus on the feedback given by the audience about their perception of the AI, as they observed it during the excerpts of the piece we performed for them. It is important to note that, before the performance, we purposely avoided informing the audience about the presence of an algorithm driving lights and music.

The two most significant comments of the audience were that some spectators did not realize that music and light were driven by an AI algorithm, and that another spectator suggested us to create more explicit interactions between the algorithm and the performers on stage. These comments are perhaps unsurprising if we consider the issue of implicit/explicit interaction discussed above. Nevertheless, they confirmed our early intuition about the challenges of making an AI perceivable to a public audience. Our algorithm manifests itself only indirectly, through lights and music and it lacks a physical embodiment. This was a choice we made in order to reinforce our critical stand on the impact of AI in society; the AIs regulating social media, welfare and warfare are pervasive and ubiquitous, while simultaneously invisible and implicit, integrated seamlessly as they are in institutional structures and interpersonal lives. Our AI was designed with this in mind, and is therefore omnipresent and unobservable.

However, concept aside, on an actual theater stage it is indeed a challenge to manifest the AI's presence and expressivity while completely avoiding a physical representation of it (and potentially fall into naive anthropomorphism). Thus, while we still believe the concept we explored is meaningful to the overall piece, the audience feedback made us realize that, in order to be perceived as a real presence, the AI in *Humane Methods* needs to be embodied and its agency has to be manifested, in a way or another. This is still ongoing research, thus we can only speculate here, but we believe this issue opens a fascinating challenge regarding the fragile balance between AI in control versus humans in control elaborated in the previous section.

## 5 Discussion

This chapter focused on possible methodological approaches to the use of ML and AI in the context of music and performance creation. We have offered an autoethnographic perspective on a research and creation process spanning five years. Our aim was to open

up our process and proffer insight into one possible way of coupling art and science in the application of AI and ML. In closing this chapter, we discuss three main aspects of the research. The first relates to how we approached AI and ML in relation to music and sound creation. The second concerns a shift in terminology that accompanied, and became manifested through, our projects. The third and final one addresses the hybrid methodological approach we developed progressively and contextually to each artwork.

## 5.1 Artificial Intelligence and Music

Many music-related fields are currently facing important changes due to the intervention of machine learning and artificial intelligence technology in music processing. These changes occur at different levels: creation, consumption, production and diffusion [13]. While AI technologies offer increasingly broad options to researchers and artists alike, with the present contribution we stress the importance of acknowledging the historicity of these technologies where and how they emerged and developed as well as the increasingly complex implications of their interaction with multiple layers of humans' societies. As we have demonstrated through the review and discussion of our specific collaborative projects, the use of ML and AI in music-related research does not have to be constrained by a paradigm of control of humans over algorithms, of humans over musical forms, or of humans over other humans. Rather, by nurturing in-depth and long-term art and science collaborations it is possible to define new modalities of interaction across ML, scientific research and artistic creation.

At a musical level, both *Corpus Nil* and *Humane Methods* foster improvisational and unconventional types of interaction between algorithms, human bodies and musical forms. What should be emphasised, however, is that these kinds of musical interactions emerged from the combination of particular scientific and artistic drives which, in turn, evolved out of specific socio-cultural contexts. Artistic intervention into scientific research was coupled with practice-based research in ways that ramified our thoughts and strategies in multiple directions. Our own idea of what computational music can be was heavily affected by the particular paths we chose throughout the past five years of collaborative work. Whereas this allowed the germination of nuanced and technically rigorous interaction systems, it also made the development of satisfying outcomes slower than it could have been, had either of chosen to work alone. Thus, the same combination of practices and approaches which affords us with the capacity to deeply couple art and science, is a factor that impedes fast turnarounds and requires mid-large timescales in order to reveal artistically and scientifically valid outcomes.

## 5.2 From Machine Learning to Artificial Intelligence

Our work witnessed a change in terminology that has been occurring since the deep learning breakthrough in 2012, that is, the shift from “machine learning” to “artificial intelligence”. Machine learning used to – and, to some extent, still does – refer to the technological underpinnings that allow computational systems to “learn” from data, where learning means extracting structural components from a dataset. The term artificial intelligence is however more ambiguous: it was originally coined in the 1950’s to designate a symbolic approach for modeling human cognition, but nowadays what is called AI mostly relies on vector-space generated by deep neural networks (which is not a symbolic approach) [16]. The term has its advantage to communicate about the technology and to create research and business incentives. Then, we also believe that the term imposed itself following the recent capacities of machine learning techniques. When AlphaGo [52] won against the worlds best Go player, the idiom “machine learning” did not seem comprehensive enough to designate the technology that made it possible, its computational capacities, or the new prospect that such technology had opened up.

While, today, the HCI community often uses the term Artificial Intelligence, in our own research we preferably adopt the term Machine Learning to designate the type of algorithms we develop and use. Nevertheless, in our view, the adoption of the term AI triggers interesting questions into the concepts that are most commonly associated with the term. Speaking of AI triggers a specific and ambiguous imaginary that is worth understanding in detail in order to make better use (or no use at all) of this technology as practitioners and researchers in HCI.

The dialectical shift can be traced in our own work. Our first project, *Corpus Nil*, relied on a low-capacity movement recognition system. In this case, we were certain that Machine Learning was the right designation because the algorithm did not manifest a distinctive behavior. Rather, its role was to recognize and learn given aspects of a performer’s physical expressivity, so as to emphasize elements of human corporeality. On the contrary, the second project, *Humane Methods*, involved a radically different algorithm chosen not for its capacity to achieve a certain goal, but rather for the particular computational behavior driving its learning mechanism. In this case, the term Machine Learning felt constraining, for it precluded an understanding of the algorithm as a socio-cultural phenomenon, beyond its computational capabilities. The algorithm had almost become a black box, interpretable only by observing its choices and subsequent actions. In this sense, it became observable as an actor.

### 5.3 Hybrid Methodology

By observing the varied methods we deployed in the research presented in this chapter, what comes to the fore is a hybrid methodology, an adaptive range of methods which evolved according to the nature and context of the particular problem at hand. This is illustrated by how, in *Corpus Nil* and *Humane Methods*, learning algorithms were used in two fundamentally different ways, and how, in each case, the coupling of the respective needs of art and science led to unconventional and effective results.

In *Corpus Nil*, we used an algorithmic system to design expressive, body-based musical interactions between a human performer and algorithms. We set out to explore Machine Learning algorithms as a *tool*, therefore our focus was less on the algorithm itself and more on what results the algorithm could achieve during the performance. The particular choreography created for the piece revealed limits of available computational approaches for the design of interactions with sound, and led us to find an ad hoc method based on physiological sensing. This method – triggered by artistic intervention into scientific research – yielded a new approach to the analysis of gesture expressivity, one based on complimentary aspects of muscle physiological data. This allowed us to capture dimensions of expressivity that greatly differ from those obtained through computer vision interfaces, which are the most commonly used in the literature.

In *Humane Methods*, we designed a computational framework which eventually became an actor, an independent entity directing sound and light in a dance-theatre performance according to its own internal functioning. Contrarily to *Corpus Nil*, at the core of this piece was an investigation of AI as a cultural and social *concept*, therefore we were less concerned with the results the algorithm could achieve and more preoccupied with how to manifest the functioning of the algorithm itself and to exhibit its behavior. On one hand, this shift of focus proved challenging because it forced our research to confront a broad set of (scientific, artistic and popular) expectations surrounding the notion of AI. On the other hand, though, the challenge has opened before us a vast field of investigation, which extends across HCI, performing art and into the socio-political impact of AI. We feel that the prospect for a further entanglement of science, art and politics is, today, perhaps more urgent than ever.

## 6 Acknowledgements

We would like to thank Frédéric Bevilacqua (IRCAM, Sorbonne Université), Sarah Fdili Alaoui and Margherita Pevere for their helpful comments on the paper. This research was partly supported by the European Research Council under the European Union’s Seventh

Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 283771; by the CNRS under the programme PEPS / project INTACT; and by La Diagonale Paris-Saclay, under the programme Phares 2019 / project Humane Methods.

## References

- [1] S. Amershi, M. Cakmak, W. B. Knox, and T. Kulesza. Power to the people: The role of humans in interactive machine learning. *Ai Magazine*, 35(4):105–120, 2014.
- [2] S. Amershi, D. Weld, M. Vorvoreanu, A. Fourney, B. Nushi, P. Collisson, J. Suh, S. Iqbal, P. N. Bennett, K. Inkpen, et al. Guidelines for human-ai interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2019.
- [3] J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine bias. *ProPublica*, May, 23:2016.
- [4] J. Bardzell, S. Bardzell, and L. Koefoed Hansen. Immodest proposals: Research through design and knowledge. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 2093–2102, 2015.
- [5] S. Benford, C. Greenhalgh, G. Giannachi, B. Walker, J. Marshall, and T. Rodden. Uncomfortable interactions. In *Proceedings of the sigchi conference on human factors in computing systems*, pages 2005–2014, 2012.
- [6] M. Biggs. Learning from experience: approaches to the experiential component of practice-based research. *Forskning-Reflektion-Utveckling. 6-21. Stockholm: Swedish Research Council, Vetenskapsr det*, 2004.
- [7] L. Blackman. *Immaterial Bodies : Affect, Embodiment, Mediation*. Sage Publications, London, 2012.
- [8] M. Blythe, K. Andersen, R. Clarke, and P. Wright. Anti-solutionist strategies: Seriously silly design fiction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4968–4978, 2016.
- [9] S. Bødker. When second wave hci meets third wave challenges. In *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*, pages 1–8, 2006.
- [10] C. J. Cai, S. Winter, D. Steiner, L. Wilcox, and M. Terry. ”hello ai”: Uncovering the onboarding needs of medical practitioners for human-ai collaborative decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–24, 2019.
- [11] L. Candy. Practice based research: A guide. *CCS report*, 1:1–19, 2006.
- [12] B. Caramiaux, M. Donnarumma, and A. Tanaka. Understanding gesture expressivity through muscle sensing. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 21(6):1–26, 2015.

- [13] B. Caramiaux, F. Lotte, J. Geurts, G. Amato, M. Behrmann, F. Bimbot, F. Falchi, A. Garcia, J. Gibert, G. Gravier, et al. Ai in the media and creative industries. *[Research Report] New European Media (NEM)*, pages 1–35, 2019.
- [14] B. Caramiaux, N. Montecchio, A. Tanaka, and F. Bevilacqua. Adaptive gesture recognition with variation estimation for interactive systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 4(4):1–34, 2015.
- [15] B. Caramiaux and A. Tanaka. Machine learning of musical gestures: Principles and review. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME)*, pages 513–518. Graduate School of Culture Technology, KAIST, 2013.
- [16] D. Cardon, J.-P. Cointet, A. Mazières, and E. Libbrecht. Neurons spike back. *Réseaux*, (5):173–220, 2018.
- [17] K. Crawford and T. Paglen. Excavating ai: The politics of images in machine learning training sets, 2019.
- [18] M. Donnarumma. XTH Sense: A Study of Muscle Sounds for an Experimental Paradigm of Musical Performance. In *Proceedings of the International Computer Music Conference*, Huddersfield, 2011. Maynooth University.
- [19] M. Donnarumma. Biophysical Music Sound and Video Anthology. *Computer Music journal*, 39(4):132–138, 2015.
- [20] M. Donnarumma. Beyond the cyborg: performance, attunement and autonomous computation. *International Journal of Performance Arts and Digital Media*, 13(2):105–119, 2017.
- [21] M. Donnarumma. On Biophysical Music. In E. R. Miranda, editor, *Guide to Unconventional Computing for Music*, chapter 3, pages 63–83. Springer, London, 2017.
- [22] M. Donnarumma, B. Caramiaux, and A. Tanaka. Muscular interactions combining emg and mmg sensing for musical practice. In *NIME*. KAIST, 2013.
- [23] P. Dourish. *Where the action is: the foundations of embodied interaction*. MIT press, 2004.
- [24] G. Dove, K. Halskov, J. Forlizzi, and J. Zimmerman. Ux design innovation: Challenges for working with machine learning as a design material. In *Proceedings of the 2017 chi conference on human factors in computing systems*, pages 278–288, 2017.
- [25] M. D. D’Urso. Artist’s website. <http://www.disorienta.org/en/>.
- [26] E. Edmonds and L. Candy. Relating theory, practice and evaluation in practitioner research. *Leonardo*, 43(5):470–476, 2010.
- [27] J. A. Fails and D. R. Olsen Jr. Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces*, pages 39–45, 2003.



- [28] S. Fdili Alaoui. Making an interactive dance piece: Tensions in integrating technology in art. In *Proceedings of the 2019 on Designing Interactive Systems Conference*, pages 1195–1208, 2019.
- [29] R. Fiebrink and B. Caramiaux. The machine learning algorithm as creative musical tool. In *The Oxford Handbook of Algorithmic Music*. Oxford University Press, 2018.
- [30] R. Fiebrink, P. R. Cook, and D. Trueman. Human model evaluation in interactive supervised learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 147–156, 2011.
- [31] J. Françoise. *Motion-sound mapping by demonstration*. PhD thesis, Université Pierre et Marie Curie, Paris, 2015.
- [32] J. Françoise, N. Schnell, R. Borghesi, and F. Bevilacqua. Probabilistic models for designing motion and sound relationships. 2014.
- [33] M. Fuller and O. Goriunova. *Bleak Joys*. University of Minnesota Press, Minneapolis, MN, 2019.
- [34] N. Gillian and J. A. Paradiso. The gesture recognition toolkit. *The Journal of Machine Learning Research*, 15(1):3483–3487, 2014.
- [35] M. Gillies. Understanding the role of interactive machine learning in movement interaction design. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(1):1–34, 2019.
- [36] M. Gillies, R. Fiebrink, A. Tanaka, J. Garcia, F. Bevilacqua, A. Heloir, F. Nunnari, W. Mackay, S. Amershi, B. Lee, et al. Human-centred machine learning. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 3558–3565, 2016.
- [37] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal processing magazine*, 29(6):82–97, 2012.
- [38] K. Höök. Transferring qualities from horseback riding to design. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*, pages 226–235, 2010.
- [39] H. Hutchinson, W. Mackay, B. Westerlund, B. B. Bederson, A. Druin, C. Plaisant, M. Beaudouin-Lafon, S. Conversy, H. Evans, H. Hansen, et al. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 17–24, 2003.
- [40] L. Irani, J. Vertesi, P. Dourish, K. Philip, and R. E. Grinter. Postcolonial computing: a lens on design and development. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1311–1320, 2010.

- [41] S. Katan, M. Grierson, and R. Fiebrink. Using interactive machine learning to support interface development through workshops with disabled people. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 251–254, 2015.
- [42] O. Keyes, J. Hoy, and M. Drouhard. Human-computer insurrection: Notes on an anarchist hci. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2019.
- [43] R. B. Knapp and H. S. Lusted. A Bioelectric Controller for Computer Music Applications. *Computer Music Journal*, 14(1):42–47, apr 1990.
- [44] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [45] O. Kulesz. Culture, platforms and machines: the impact of artificial intelligence on the diversity of cultural expressions. *Intergovernmental Committee for the Protection and Promotion of the Diversity of Cultural Expressions; Paris*, 2018.
- [46] R. Malina. Welcoming Uncertainty: The strong case for coupling the contemporary arts to science and technology. In J. Scott, editor, *Artists-in-labs : Process of Inquiry*, page 15. Springer, Wien, 2006.
- [47] K. Murobushi. Artist’s website. <https://ko-murobushi.com/eng/>.
- [48] M. Pevere. Artist’s website. <https://margheritapevere.com>.
- [49] D. Rosenboom, L. M. Number, and D. Rosenboom. *Extended Musical Interface with the Human Nervous System*, volume 1990-97. Leonardo, 1990.
- [50] H. Scurto, B. Van Kerrebroeck, B. Caramiaux, and F. Bevilacqua. Designing deep reinforcement learning for human parameter exploration. *arXiv preprint arXiv:1907.00824*, 2019.
- [51] M. Shildrick. *Embodying the Monster : Encounters with the Vulnerable Self*. Sage Publications, London, 2002.
- [52] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- [53] I. Stengers. *The Invention of Modern Science*. University of Minnesota Press, Minneapolis, MN, 2000.
- [54] A. Tanaka. Musical technical issues in using interactive instrument technology with application to the BioMuse. In *Proceedings of the International Computer Music Conference*, pages 124–126, 1993.
- [55] M. Ware, E. Frank, G. Holmes, M. Hall, and I. H. Witten. Interactive machine learning: letting users build classifiers. *International Journal of Human-Computer Studies*, 55(3):281–292, 2001.

- [56] Q. Yang, A. Steinfeld, and J. Zimmerman. Re-examining whether, why, and how human-ai interaction is uniquely difficult to design. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020.