Machine Learning for Computer Music Multidisciplinary Research: A Practical Case Study

Hugo Scurto, Axel Chemla–Romeu-Santos

Machine Learning for Computer Music Multidisciplinary Research: A Practical Case Study

Hugo Scurto*[1] & Axel Chemla–Romeu-Santos*[2,1]

[1] STMS IRCAM–CNRS–Sorbonne Université
[2] Laboratorio d'Informatica Musicale, Università degli Studi di Milano
{scurto,chemla}@ircam.fr

Abstract. This paper presents a multidisciplinary case study of practice with machine learning for computer music. It builds on the scientific study of two machine learning models respectively developed for data-driven sound synthesis and interactive exploration. It details how the learning capabilities of the two models were leveraged to design and implement a musical instrument focused on embodied musical interaction. It then describes how this instrument was employed and applied to the composition and performance of *ægo*, an improvisational piece with interactive sound and image for one performer. We discuss the outputs of our research and creation process, and build on this to expose our personal insights and reflections on the multidisciplinary opportunities framed by machine learning for computer music.

Keywords: Multidisciplinarity, Machine Learning, Interface Design, Composition, Performance

1 Introduction

Machine learning is a field of computer science that studies statistical models able to automatically extract information from data. The statistical learning abilities of the models induced a paradigm shift in computer science, which reconsidered mechanistic, rule-based models, to include probabilistic, data-driven models. Recent applications of machine learning led to critical advances in disciplinary fields as diverse as robotics, biology, or human-computer interaction. It also contributed to new societal representations of computers through the loosely-defined notion of Artificial Intelligence (AI).

Computer music also witnessed an increased interest in machine learning. Research has mostly been scientific in focus, using and studying models to automatically analyse musical data—*e.g.,* extracting symbolic information related to pitch or timbre from audio data. This led to technical advances in the field of music information retrieval [1], while also benefiting the field of musicology, notably through large-scale computational analysis [2]. In parallel, machine learning also

* Equal contribution.

enabled the building of many automatic music generation systems, which are currently being invested by the industry in the wave of AI [3].

Importantly, these scientific investigations of machine learning have also enabled the birth of new musical practices. For example, gesture modelling, as a scientific challenge, opened new design perspectives on body-based musical instruments that adapts to one's way of playing it [4]. Similarly, symbolic sequence modelling created new human-machine improvisational situations where the machine learns to imitate a musician's style [5]. Reciprocally, artistic investigations of machine learning began taking a complementary approach, using the models themselves as material for composition of sound [6] and image [7].

We are interested in adopting a *joint scientific and musical approach* to machine learning research. We are inspired by the computer music pioneer Jean-Claude Risset [8], whose research and creation approach to computer science enabled new scientific understandings of sound as a physical and perceptual phenomenon, jointly with an artistic commitment toward the computed aesthetics. His work and personal approach gave insight to both scientists—ranging from formal to social science—, and artists—ranging from composers and performers to instrument designers. Our wish is to perpetuate his multidisciplinary impetus toward contemporary computer music issues related to machine learning.

The work that we present here is a step toward this direction. We led a *scientific* investigation of two machine learning models that jointly frame new data-driven approaches to sound synthesis. We then adopted a *musical* approach toward these models, leveraging their interactive learning abilities to design a musical instrument, which we employed to create an improvisational piece. Rather than seeking general abstractions or universal concepts, our wish was to test these models through a practical case study to engage a personal reflection on the musical representations and behaviors that they may encode. Our hope is that our idiosyncratic research and creation process will help open multidisciplinary perspectives on machine learning for computer music.

The paper is structured as follows. We start by the scientific foundations of our work, describing the two models that we developed for two musical issues—sound analysis-synthesis, and sonic exploration. Next, we present the design of our musical instrument, by describing its workflow and implementation with a focus on embodied musical interaction. We then describe *ægo*, an improvisational piece with interactive sound and image for one performer, which we wrote for our instrument. Finally, we discuss our research and creation process to draw conceptual insight on machine learning for computer music from crossed science, design, and art perspectives.

2 Scientific Modelling

In this section, we describe our two machine learning models, based on *unsupervised learning* and *reinforcement learning*, from a computer science perspective. We explain how they respectively address two specific musical issues: sound synthesis-analysis and sonic exploration.

2.1 Unsupervised Learning for Sound Analysis and Synthesis

Musical Issue. Most sound analysis-synthesis techniques, such as the phase vocoder [9] or the wavelet transform [10], are based on invertible transforms that are independent of the analyzed sounds. Such transforms provide frameworks that can be applied regardless to the nature of the signal, but in return impose a determined structure such that the extracted features are not corpus-dependant. Conversely, could we think about a method retrieving continuous parameters from a given set of sounds, but rather aiming to recover its underlying structure?

Model. The recent rise of *unsupervised generative models* can provide a new approach to sound analysis-synthesis, by considering each item of a given audio dataset $\{\mathbf{x}_n\}_{n\in 1\dots D}$, in our case a collection of spectral frames, as draws from an underlying probability distribution $p(\mathbf{x})$ that we aim to recover. The introduction of latent variables $\mathbf{z}$ allows us to control a *synthesis* process by modelling the joint distribution $p(\mathbf{x},\mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})$, such that these variable act as parameters for the generative process $p(\mathbf{x}|\mathbf{z})$. The full inference process, that would here correspond to the *analysis* part, leverages the Bayes' rule $p(\mathbf{z}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{p(\mathbf{x})}$ to recover the distribution $p(\mathbf{z}|\mathbf{x})$, called the posterior.

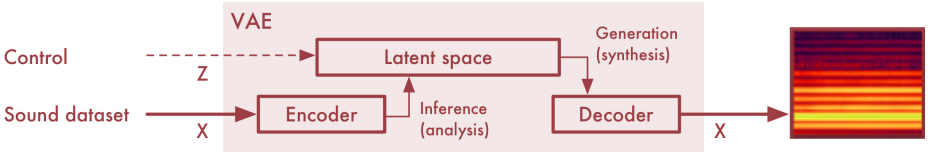

Fig. 1. Unsupervised learning for sound analysis and synthesis. The variational auto-encoder (VAE) encodes a sound dataset into a high-dimensional latent space, which can be parametrically controlled to synthesize new sounds through a decoder.

To improve expressivity of inference and generation, we propose to investigate variational learning, a framework approximating the true posterior $p(\mathbf{z}|\mathbf{x})$ by a distribution $q(\mathbf{z}|\mathbf{x})$, such that both inference and generative process can be freely and separately designed, with arbitrary complexity. The variational auto-encoder (VAE) is representative of such methods [11]. In this model (Fig. 1), inference and generation processes are held by two jointly trained separated networks, respectively the *encoder* and the *decoder*, each modelling respectively the distributions $q(\mathbf{z}|\mathbf{x})$ and $p(\mathbf{x}|\mathbf{z})$. The inherent Bayesian nature of variational learning forces the smoothness of the *latent space*, a high-dimensional, non-linear sonic space, whose parametric dimensions can be freely explored in the manner of a synthesizer.

In related work, we show how this latent space can be regularized according to different criteria, such as enforcing perceptual constraints related to timbre [12]. We refer the reader to the latter paper for technical details on the model and quantitative evaluation on standard sound spectrum datasets.

2.2 Reinforcement Learning for Sonic Exploration

Musical Issue. Sonic exploration is a central task in music creation [13]. Specifically, exploration of digital sound synthesis consists in taking multiple steps and iterative actions through a large number of technical parameters to move from an initial idea to a final outcome. Yet, the mutually-dependent technical functions of parameters, as well as the exponential number of combinations, often hinder interaction with the underlying sound space. Could we imagine a tool that would help musicians explore high-dimensional parameter spaces?

Model. We propose to investigate *reinforcement learning* to support exploration of large sound synthesis spaces. Reinforcement learning defines a statistical framework for the interaction between a learning agent and its environment [14]. The agent can learn how to act in its environment by iteratively receiving some representation of the environment's state S, taking an action A on it, and receiving a numerical reward R. The agent's goal, roughly speaking, is to maximize the cumulative amount of reward that it will receive from its environment.

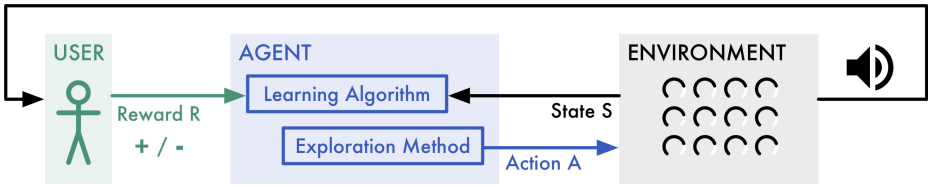

Fig. 2. Reinforcement learning for sonic exploration. The agent learns which actions to take on a sound synthesis environment based on reward given by the musician. The agent implements an exploration method to foster discovery along interaction.

For our case of sonic exploration, we propose that the musician would listen to the agent exploring the space, and teach it how to explore by giving reward data (Fig. 2). Formally, the environment's state is constituted by the numerical values of all synthesis parameters. The agent's actions are to move one of the parameters up or down at constant frequency. Finally, the musician communicates *positive or negative reward* to the agent as a subjective feedback to agent actions. We implemented a deep reinforcement learning model to support learning from human reward signal in high-dimensional parametric spaces [15].

A crucial requirement for reinforcement learning agents is to *autonomously explore their environment*, to keep on discovering which actions would yield the most reward. We developed a statistical method, based on intrinsic motivation, which pushes the agent to "explore what surprises it". The resulting interactive learning workflow was found to be useful to relax musicians' control over all synthesis parameters, while also provoking discoveries by exploring uncharted parts of the sound space. We report the reader to [16, 17] for technical details on the tool and qualitative evaluation from expert sound designers.

3 Instrument Design

In this section, we present our musical instrument that combines our two models and leverages their learning capabilities from a design perspective. We describe how interaction design was framed in joint coordination with hardware and software engineering to support embodied musical interaction.

3.1 Interaction design

Motivation. Our main design motivation was to use our reinforcement learning agent to support musical exploration of high-dimensional latent sound spaces built by our unsupervised learning model.

Specifically, our aim was to exploit the exploration behaviour of our reinforcement learning agent to support non-symbolic *improvisation* inside the spaces. Instead of acting as a tool, we used machine learning as an expressive partner [5] that would be playable by musicians using positive or negative feedback.

A complementary aim was to employ the generative abilities of our unsupervised learning model to support *customization* of sound synthesis spaces. Instead of accurately modelling sounds, we used machine learning as a creative interface [18] that lets musicians experiment with the nonlinearities of the latent spaces.

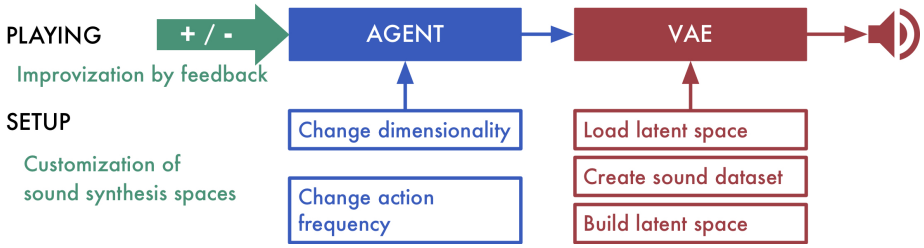

Fig. 3. The interactive workflow that we designed for our instrument.

Workflow. We designed a two-phase interactive workflow, shown in Fig. 3.

The *setup* phase allows musicians to configure the instrument. They can create a customized sound dataset for the unsupervised learning model, experiment with various training parameters, or also load a previously-built latent sound space. They can also change dimensionality of the reinforcement learning agent to explore specific dimensions of the latent sound space, as well as the frequency at which it would take actions inside the latent space.

The *playing* phase allows musicians to improvise with the agent by means of feedback. The agent produces a continuous layer of sound from the spectrum output of the VAE. Musicians can either cooperate with its learning by giving consistent feedback data to attain a sonic goal. Or, they can obstruct its learning by giving inconsistent feedback data to improvise through sonic exploration.

3.2 Engineering

Implementation. Technically (see Fig. 4), the reinforcement learning agent receives a representation of the environment's state S as a position in the latent space $\mathbf{z}$. Then, it takes an action A corresponding to a displacement along some dimension of the latent space. The resulting position has the unsupervised learning model generate a sound spectrum $\mathbf{x}$. Based on the sound, the musician would communicate reward R to the agent. The latter would progressively learn to explore the latent space in relation to the musician's feedback data.

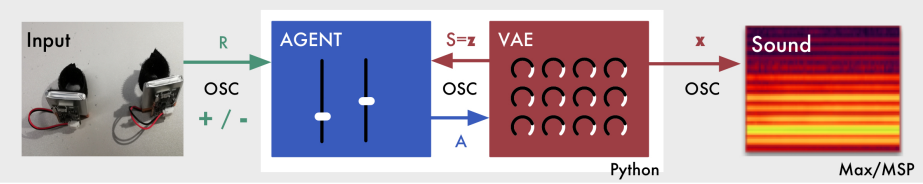

Fig. 4. Schematic representation for the engineering of our instrument.

Hardware. We designed a hardware prototype to support embodied musical interaction (see Fig. 4, left). It consists in two velcro rings, each of them equipped with a wireless inertial measurement unit[1]. We took each unit angular rotation about each forearm axis and summed them to compute a single, normalized numerical reward signal. This, combined with the lightweight, nonintrusive velcro rings, lets musicians experiment with a wide range of gesture vocabulary [19] to communicate positive or negative feedback to the agent.

Software. We implemented our two machine learning models as Python libraries[2,3]. We developed a Max/MSP patch to implement a user interface for the setup phase, as well as a hardware data converter for the playing phase. We leveraged the OSC protocol to bridge hardware data, reinforcement learning agent, unsupervised latent space, and sound spectra together into the patch.

4 Musical Artwork

In this section, we present *ægo*, an improvisational piece that we wrote for our musical instrument, to be premiered at the *14th International Symposium on Computer Music Multidisciplinary Research*, held in Marseille, France. We describe the intended aesthetics of sound, image and body, and detail how composition and performance were approached in relation to our learning instrument.

[1] `http://ismm.ircam.fr/riot/`
[2] `https://github.com/acids-ircam/variational-timbre`
[3] `https://github.com/Ircam-RnD/coexplorer`

4.1 Aesthetics

Motivation. Our artistic motivation for *ægo* was to open a sensitive reflection on what may actually be learned on a musical level through interaction with machine learning, both by the human and its artificial alter ego—the machine. To share this reflection with members of an audience, we opted for a performance format that displays a human and a machine mutually learning to interact with each other—on an embodied level for the human, and on a computational level for the machine—through live improvization.

The learning machine possesses a distinctive musical behaviour, as well as two latent sound spaces, that are all originally unknown to the human performer. The latter will expressively negotiate control of these spaces with the machine, communicating positive or negative feedback using our instrument and its motion sensors placed in both hands. The slowly-evolving spectromorphologies, synthesized and projected in real-time on stage, create a contemplative, minimalist atmosphere intended to let members of the audience freely consider potential learnings of musical qualities by the human and the machine.

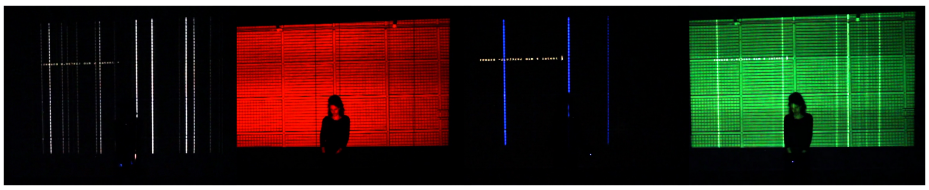

Fig. 5. Pictures taken from *ægo*.

Intentions. The piece's aesthetic intentions toward machine learning lie at three intertwined levels: sound, image, and body (see Fig. 5).

One of our intentions was to reveal the *sound representations* learned by the unsupervised learning model to the audience. We thus built latent sound spaces using sound data that was commonly used and produced in pioneering works of computer music. In addition, we projected the generated sound spectrums on stage to provide the audience with a visual representation that accentuate, not disrupt, the sonic perception of the piece.

Another intention was to display the *exploration behaviour* of the reinforcement learning agent in front of the audience. To do this, we wanted to challenge the skills and abilities usually at stake in performance, by summoning an ecological approach and evoking a sense of reciprocal interaction between the human and the machine. In this sense, rather than using it for control purposes, we used the body of the performer to convey kinesthetic information about how machine exploration may be internally experienced by a human. In parallel, we added raw textual information about the machine's internal state at top left of the image projection to emphasize the machine's encoded perception of the performer.

4.2 Writing

Composition. The piece was composed at three temporal scales (see Fig. 6).

The first scale is that of *exploration*. It consists in the improvisational paths taken by the reinforcement learning agent following the performer's feedback data. We set the frequency of agent actions between 30 and 100 milliseconds. This choice allowed for slow, continuous evolution of spectromorphologies, which enables to grasp the behaviour of the agent inside the latent spaces.

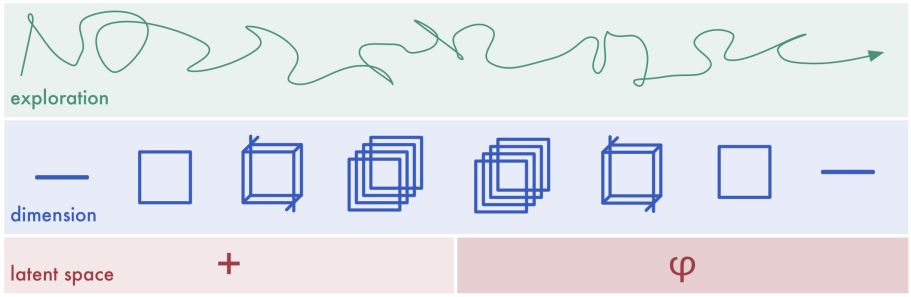

Fig. 6. Temporal structure composed for the piece.

The second scale is that of latent space *dimensionality*. It consists in defining the axis of the latent spaces that the reinforcement learning agent will explore. We set the dimensions to 1, 2, 4, and 8, respectively. This allows to write a specific kind of musical form inside the latent space: the more dimensions we open to the agent, the more sonic variance the performer and audience members will experience.

The third scale is that of latent space itself. It consists in connecting the reinforcement learning agent to another type of latent space. We used two latent spaces, respectively built from additive synthesis sounds and physical instruments recordings (flute, saxophone, piano, violin, bassoon [20]). This enables to write form within different soundscapes, allowing the building of a narrative (here, going from elementary sinusoidal spectra to richer instrumental timbres).

Performance. While the piece is intended to be improvised, our sole instruction toward the stage performer is that he or she globally performs with the machine with an overall sense of attentiveness[4]. We propose that the performer would start the piece facing the audience, relaxed, using the instrument with small forearm rotations only. As the piece would unfold over time, the performer would be free to adapt its gestures in response to the slowly evolving complexity of the explored spaces, focusing on embodied interaction with the machine.

A second contributor is required to manage the two remaining temporal scales of the piece—*i.e.,* changing dimensionalities, and switching latent spaces.

[4]See the following video recording: `https://youtu.be/gCzOoNChlJQ`

5 Discussion

In this section, we take a critical look at the output of our case study by discussing our research and creation process. We then expose our personal reflections emerging from practice with machine learning, and propose conceptual insight for future multidisciplinary inquiries in the realm of computer music.

5.1 Case study

Process. The work presented here relates a practical case study with machine learning in the frame of computer music. We leveraged both conceptual and technical aspects of machine learning to jointly produce *scientific knowledge* with our two models for sound synthesis, as well as *musical creations* through the design of our instrument and the writing of our musical piece. In this sense, our work emerged from a research and creation process, in which we closely articulated a research methodology with a creation project.

We followed a sequential disciplinary agenda (see Fig. 7, solid lines and arrows). We started by the scientific modelling of sonic exploration and sound synthesis, which took us two years to date. We then planned a one-month period to conceive the instrument, write and practice the musical piece. This research and creation agenda was mainly required by our work occupation focusing on computer science research without necessarily addressing music creation.

While many researchers of our laboratory were involved in scientific modelling, we (the two coauthors) managed instrument design and musical piece as a pair. Importantly, we both followed a dual training in science and music, and were doctoral students in the domain of machine learning applied to computer music at the time of writing. In addition, both of us have professional experience in music composition and performance. These dual skills were central to individually work, as well as to effectively collaborate, on conceptual and technical aspects related to machine learning throughout the process.

Output. The relatively short period dedicated to musical creation pushed us to take pragmatic decisions about the form of outputs, notably by relinquishing certain technical developments. For example, using the unsupervised learning model to learn temporal features of sound spectrums could have improved the timbre richness of the generated sounds, as well as supported other musical forms than slow spectromorphology evolution. Also, other agent commands than feedback data could have been designed to support expressive human control over the reinforcement learning agent exploration. Finally, many other musical forms could have been conceived, using other sound datasets—*e.g.,* voice corpora or environmental sounds—and investigating other temporal writings for dimensionality and exploration. Future continuation of our work may consider addressing these research questions to evolve the generated outputs.

5.2 Authors' reflections on machine learning for computer music

Conceptual insight. Beyond the created outputs, our process of practice with the two machine learning models let us reflect on conceptual issues, which feed back into many different disciplines (see Fig. 7, dashed lines and arrows).

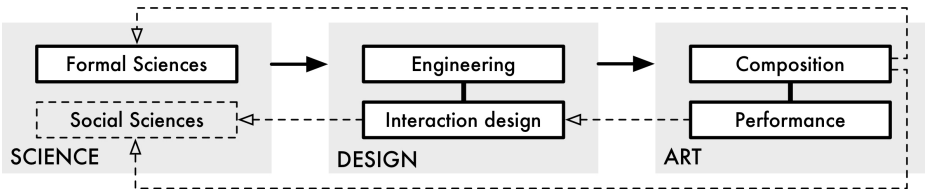

Fig. 7. Our case study. Solid arrows: The sequential research and creation process that we took to scientifically investigate our models, and musically create our instrument and artwork. Dashed arrows: The personal conceptual insight gathered along our process.

On the one hand, composing with the sonic aesthetics produced by the unsupervised learning model let us reflect on epistemological issues that span both formal and social science (Fig. 7, upper and lower dashed arrows). Should machine learning be considered as a modelling tool for sound data, or rather as a framework for sound synthesis that remains to be crafted? Our insight leans toward the latter option. Rather than imposing deterministic rules to define a sound space [21], probabilistic methods propose heuristics that aim to inverse this methods by retrieving structure directly from the data. More specifically, Bayesian approaches filtrates the "space of everything possible" to get closer from the data structure, thus providing interesting generalization abilities in addition to structural information, from the point of view of *formal science*. Conversely, adopting an artistic approach to the learned representations also provides an alternative way of evaluating these models, completing existing machine learning-focused evaluations methods of such unsupervised learning systems. However, such evaluations have to deal with musicological approaches in the realm of the *social sciences*, and remains still an underrated field of research.

On the other hand, performing with a reinforcement learning-based musical instrument offers new design and scientific views on interactivity (Fig. 7, middle dashed arrows). How should we approach an artificial musical partner that learns to behave from our sole feedback data? Alternatively, should exploration be analysed as an expressive musical behaviour? Our insight is that the data-agnostic framework of machine learning may support the development of new modalities for human-machine interaction, which may originate from the *social sciences*. In the musical domain, machine learning may be used to enhance modes of communication that already exist between musicians. Feedback, for example, is a broad communication channel that concern all types of living or nonliving systems [22]. By *designing interactions* with machine learning that rely on feedback data, we may create more accessible musical partners and in

turn instigate analytical views on these embodied notions—as it has been the case with machine learning-based gesture modelling tools [4]. Exploration, as a performative and improvisational practice, remains to be investigated more deeply in that sense.

Toward intrinsic approaches. In this paper our approach was to study the artistic possibilities emerging from the encounter of our two models, rather than to evaluate them separately on their respective tasks. Precisely, our experience in practicing such models revealed to us two distinctive approaches: an *extrinsic* approach, where machine learning models are designed towards a specific task and used faithfully to this end—such as in music information retrieval—, and an *intrinsic* approach, where these models are exploited for themselves and taken as objects that can be explored, hacked, and manipulated—such as in gesture modelling, or improvisational systems. While the first approach has so far been the most common, as machine learning was originally created to tackle complex issues that preceding techniques fell short with, we think that the second may unfold new creative opportunities for computer music, just as Jean-Claude Risset's joint scientific and musical approach to computing did [23]. We hope that the present case study stands in favour of this argument.

While we built on our joint machine learning and music training to lead our case study, it may require more time to manage collaboration between machine learning experts and researchers, engineers, musicians, artists, musicologists, scientists, designers, or epistemologists, toward shared musical goals. We believe that multidisciplinary collaboration is key to lead intrinsic examination of machine learning, and that the latter may be crucial to go beyond suspicions and actively negotiate the place of the human artist in upcoming AI music systems.

6 Conclusion

We presented a practical case study of machine learning for computer music. We studied two machine learning models, from which we designed a musical instrument, and wrote a piece for it. We discussed the research and creation process that fostered our case study and showed the conceptual benefits in terms of feedback. Future work may include multidisciplinary collaborations to intrinsically study machine learning in the realm of computer music.

Acknowledgements

We thank Frédéric Bevilacqua, Philippe Esling, Gérard Assayag, Goffredo Haus, and Bavo Van Kerrebroeck for their broad contributions on the scientific part.

References

1. Hamel, P., Eck, D.: Learning features from music audio with deep belief networks. In 11th International Society for Music Information Retrieval Conference (ISMIR) (2010)

2. Meredith, D. (Ed.): Computational music analysis (Vol. 62). Berlin: Springer (2016)
3. Briot, J-P., Hadjeres, G., and Pachet, F.: Deep learning techniques for music generation-a survey. arXiv preprint arXiv:1709.01620 (2017).
4. Bevilacqua, F., Zamborlin, B., Sypniewski, A., Schnell, N., Gudy, F., & Rasamimanana, N.: Continuous realtime gesture following and recognition. In International gesture workshop, Springer, Berlin, Heidelberg, pp. 73-84 (2009, February).
5. Assayag, G., Bloch, G., Chemilier, M., Cont, A., & Dubnov, S.: Omax brothers: a dynamic topology of agents for improvization learning. Proceedings of the 1st ACM workshop on Audio and music computing multimedia (2006)
6. Ghisi, D. Music across music: towards a corpus-based, interactive computer-aided composition. Doctoral dissertation, Paris 6 (2017)
7. Akten, M., Fiebrink, R., Grierson, M.: Deep Meditations: Controlled navigation of latent space. Goldsmiths University of London (2018).
8. Risset, J.-C.: Fifty Years of Digital Sound for Music. In: Proceedings of the 4th Sound and Music Computing Conference (SMC) (2007)
9. Rodet, Xavier and Depalle, Philippe and Poirot, Gilles : Speech analysis and synthesis methods based on spectral envelopes and voiced/unvoiced functions. European Conference on Speech Technology (1987)
10. Kronland-Martinet, R.: The wavelet transform for analysis, synthesis, and processing of speech and music sounds. Computer Music Journal, 12(4), 11-20 (1988)
11. Kingma, D., Welling, M. : Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
12. Esling, P., Chemla–Romeu-Santos, A., Bitton, A. : Bridging audio analysis, perception and synthesis with perceptually-regularized variational timbre spaces. DAFx2018 (2018)
13. Ystad, S., Aramaki, M., & Kronland-Martinet, R.: Timbre from Sound Synthesis and High-level Control Perspectives. Springer Nature (2017)
14. Sutton, R. S., & Barto, A. G.: Reinforcement learning: An introduction. MIT press (2018)
15. Warnell, G., Waytowich, N., Lawhern, V., & Stone, P.: Deep TAMER: Interactive agent shaping in high-dimensional state spaces. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018, April).
16. Scurto, H., Bevilacqua, F., & Caramiaux, B.: Perceiving Agent Collaborative Sonic Exploration In Interactive Reinforcement Learning. In: Proceedings of the 15th Sound and Music Computing Conference (SMC) (2018).
17. Scurto, H., Van Kerrebroeck, B., Caramiaux, B., Bevilacqua, F.: Designing Deep Reinforcement Learning for Human Parameter Exploration. arXiv preprint arXiv:1907.00824 (2019)
18. Fiebrink, R., Caramiaux, B., Dean, R., McLean, A.: The machine learning algorithm as creative musical tool. Oxford University Press (2016)
19. Tanaka, A., & Donnarumma, M.: The body as musical instrument. The Oxford Handbook of Music and the Body (2018)
20. Ballet, G., Borghesi, R., Hoffmann, P., & Lvy, F.: Studio online 3.0: An internet "killer application" for remote access to Ircam sounds and processing tools. In: Journées d'Informatique Musicale (JIM) (1999)
21. Chowning, J. M.: The synthesis of complex audio spectra by means of frequency modulation. Journal of the audio engineering society, 21(7), 526-534 (1973).
22. Wiener, N.: Cybernetics or Control and Communication in the Animal and the Machine. MIT press (1965)
23. Risset, J. C., & Wessel, D. L.: Exploration of timbre by analysis and synthesis. In: The psychology of music, Academic Press, pp. 113-169 (1999)