From Performance to Installation: How Interactive Reinforcement Learning Reframes the Roles of Performers and Audiences

Danny Perreault
perreault.danny@uqam.ca
Université du Québec à Montréal
Montréal, Canada

Victor Drouin-Trempe
vdrouintrempe@cvm.qc.ca
Université du Québec à Montréal
Montréal, Canada

Vincent Cusson
cusson.v@gmail.com
Université du Québec à Montréal
Montréal, Canada

David Drouin
drouin.david@courrier.uqam.ca
Université du Québec à Montréal
Montréal, Canada

Sofian Audry
audry.sofian@uqam.ca
Université du Québec à Montréal
Montréal, Canada

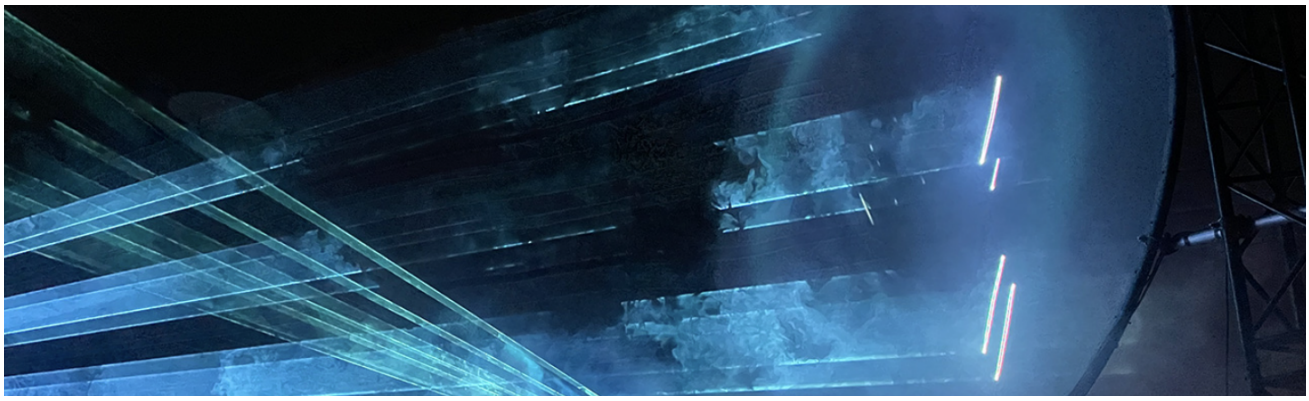

Figure 1: The *Trëma* installation, presented at Lumifest festival in 2023.

Abstract

This paper explores how interactive reinforcement learning (IRL) reconfigures the roles of performers and audiences in audiovisual performances and immersive installations. We adapt the Co-Explorer (a software tool originally developed for musical co-creation) to audiovisual immersive contexts and examine its creative potential using a reflexive research-creation approach. Our study reveals how IRL splits the role of the performer into three distinct positions: (1) the designer, who defines the parametric space; (2) the guide, who reinforces the agent's behavior; and (3) the machine performer, whose actions are shaped by interactive training. As IRL introduces agency into the creative process, it transforms traditional notions of authorship and control, enabling unexpected emergent outcomes. By showcasing an interactive installation/performance, we further explore how audiences contribute to collective creation through reinforcement-based interaction. Our findings underscore the challenges of balancing the temporality of IRL with the demands of public-facing works and of adapting RL-based systems to different exhibition contexts. Our work contributes to the discourse on co-creative systems, emphasizing the evolving roles of artists, artificial agents, and audiences in hybrid creative ecosystems.

NIME '25, June 24–27, 2025, Canberra, Australia

Keywords

Audiovisual Performance, Interactive Installation, Interactive Machine Learning, Reinforcement Learning

ACM Reference Format:
Danny Perreault, Victor Drouin-Trempe, Vincent Cusson, David Drouin, and Sofian Audry. 2025. From Performance to Installation: How Interactive Reinforcement Learning Reframes the Roles of Performers and Audiences. In *Proceedings of International Conference on New Interfaces for Musical Expression (NIME '25).* ACM, New York, NY, USA, 6 pages.

1 Introduction

Recent advances in artificial intelligence (AI) are radically transforming creative practices in nearly every artistic domain [1, 4, 20]. The integration of these technologies is reshaping the roles of artists, musicians, composers, and audiences, paving the way for new forms of human-machine co-creation [12]. Our study investigates the adaptation of the Co-Explorer[1], an interactive tool initially designed for musical co-creation, to the fields of audiovisual performance and immersive installation [17]. Applying a reflexive research-creation approach [3, 16], we examine how this transposition influences creative dynamics and transforms the roles of performers, artificial agents, and spectators.

Departing from the romantic anthropocentric ideal of solitary "creative genius", we draw on the notion of creativity as a situated phenomenon distributed across and emerging from a network of multiple agents and materials. Based on Margaret Boden's definition of creativity as the production of novelty and value [2], we adopt Oliver Bown's emphasis on creativity as a *situated*

[1] https://github.com/Ircam-RnD/coexplorer

production of novelty emerging from socio-cultural and technical networks [4]. This shifts focus from individual authorship to the distribution of *roles* between agents and their relationships within the creative process.

Through this research, we aim to expand the discourse on co-creative systems by examining an agent-based tool rooted in live human-computer interaction. We examine our collective creative process as we are using the Co-Explorer, looking at how it transforms our practice and our roles. Our methodology adopts the perspective of *creation-as-research* [7] to investigate the evolving relationship between human and machine creativity. We rely on reflexive methods such as autoethnographic vignettes [10] and group discussions to critically reflect on our experiences and interactions with the Co-Explorer without interrupting the flow of creation.

The category of AI techniques we explore, *reinforcement learning* (RL), distinguishes itself from other machine learning approaches commonly used in the arts. Supervised learning, which relies on labeled data to generate stylistically consistent outputs, has been widely adopted in audio generation systems such as WaveNet [8], while unsupervised learning approaches such as Realtime Audio Variational AutoEncoder (RAVE)[2], have supported the creation of novel sound textures and experimental compositions [5]. Unlike supervised and unsupervised learning, which operate on static datasets and focus on stylistic reproduction or pattern discovery, reinforcement learning is driven by an artificial agent that actively learns through continuous interaction with its environment. This agent makes decisions, receives feedback in the form of rewards or penalties, and adapts its behavior over time to maximize long-term outcomes, making RL particularly well suited for dynamic, real-time co-creative processes in live artistic contexts [18].

By introducing a computational adaptive agent into the creative process, RL not only transforms how practitioners collaborate with machines but also redefines audiences' experience. In our work with the Co-Explorer, the RL agent intervenes in real-time to generate audiovisual content while human practitioners train it in real-time based on their own preferences. This dynamic interaction facilitates the emergence of unexpected creative outcomes, challenging traditional notions of control and authorship. For human performers, the RL agent functions as an active collaborator, while for audiences, it adds a layer of unpredictability and autonomy. By focusing on the exploratory potential of RL, our work emphasizes the transformative possibilities of interactive co-creative machine learning systems.

In the following sections, we present three case studies showing how we adapted the Co-Explorer to different creative contexts. First, we describe our initial experiments with the Co-Explorer in a collaborative audiovisual performance setting, focusing on the dynamics between human guidance and machine autonomy. Next, we explore its deployment as an interactive public installation, examining how audience participation influences the agent's behavior and the artistic results. Finally, we discuss its integration into a gallery environment, where more controlled conditions allowed for deeper exploration of co-creative potential. These case studies are followed by a discussion that reflects on the transformative impact of interactive reinforcement learning on artistic practices, its potential for shaping new forms of human-machine collaboration, and the challenges encountered in different exhibition contexts. We conclude by considering the broader implications of our findings for the development of co-creative systems in art and music.

[2]https://github.com/acids-ircam/RAVE

2 Case Studies

2.1 First Jam with the Co-Explorer

This first case study explores the Co-Explorer's potential to reframe the dynamics of creative collaboration [11], highlighting its challenges and emergent possibilities as an interactive agent in audiovisual performances. The Co-Explorer was designed to "investigate artificial agents using deep reinforcement learning to explore parameter spaces in partnership with users for sound design" [17, p. 1]. Unlike traditional reinforcement learning systems, which optimize for a fixed objective reward function, Co-Explorer learns through subjective real-time user feedback. This approach seems particularly appropriate to artistic contexts whose subjective and exploratory nature makes optimizing for a predefined objective goal counterproductive [1]. By allowing the agent to adapt to subjective tastes without interrupting creative flow, the system fosters a unique co-creative relationship with users.

Artists and composers are used to creating sound and visual compositions by directly manipulating media material, starting from an artistic intention that they seek to implement using various compositional tools. But what happens to this creative process when the main compositional tool involves an artificial agent which we can only control indirectly? Rather than *composing* in the traditional sense, the human creator's role evolves toward that of *guiding* the agent's exploration and adaptation, reinforcing or discouraging specific outcomes based on their aesthetic preferences.

In applying this process with the Co-Explorer, we found that it was necessary to spend several hours with the agent to get interesting results. This is a notorious aspect of working with machine learning systems creatively for at least two reasons: (1) adjusting them involves providing a lot of data; and (2) since they do not behave in a rule-based manner, they somehow operate beyond rationality and need to be understood in a more embodied manner [1, pp. 63-68].

The fluctuation of our own guidance further prolonged the learning process. Because our aesthetic tastes are subjective and variable, the agent does not have to organize itself according to a precise goal, there is no *a priori* finality. Over time, the artists' desires may fluctuate, confusing the agent by sending contradictory feedback, thus potentially hampering the learning process. As an example, during our first interactions with the agent, we tended to give positive feedback when the agent led us to unusual places, creating surprising audiovisual forms. However, over time, a certain fatigue developed from observing the same audiovisual patterns over and over, leading us to attribute negative feedback to parameters which we initially rewarded positively. Rather than seeking optimization over an absolute and stable behavior, the agent thus needs to make choices that reflect unstable or contradictory rewards.

Another important realization was that, as the artificial agent moves in a multi-dimensional parametric space, it is necessary to determine in advance which audiovisual parameters will be controlled. A certain framework must therefore be specified from the outset. We have found, after experimentation, that this space of possibility must be neither too large nor too small.

These initial experiments with the Co-Explorer have shown us some capabilities and limits of interaction between an artist and an artificial agent generated by RL in creative endeavors.

Figure 2: The two screens installed above the public.

The absence of a predefined purpose in this process highlights the fluidity and adaptability required, with artists adjusting their feedback to guide the agent without rigidly dictating outcomes. This path requires artists to strike a careful balance in defining the framework within which the artificial agent operates. Crafting the parametric space to implement a coherent artistic vision while remaining open to unexpected innovations is a crucial component of the work.

One of the more surprising developments was the gradual reduction in our need to guide the agent. Over time, as its behavior stabilized through extended interactive training, we found ourselves intervening less and less. Eventually, the system no longer required active reinforcement and would navigate the parameter space autonomously. With this shift, our role subtly transformed from active performers to system designers, and finally to passive observers. Alongside this transition, the nature of the work itself evolved. What began as a live performance project had transformed into a generative installation.

As artists, we found the *process* of guiding the Co-Explorer to be an engaging interactive experience. What would happen if, instead of being shaped by our preferences as artists, the agent was influenced by audience members? The result would no longer be dictated by a unique artistic vision, but would rather evolve from the accumulated feedback of several people over time, embodying the aesthetic preferences of a collective. Thus, the idea of *Trëma* was born.

2.2 Trëma: Installation in a Public Setting

The project was showcased at two events, both named Lumifest, in the cities of Repentigny and Longueuil in Canada. These events took place in public spaces and featured various luminous digital scenographic artworks [15]. As we can see in Figure 2, the *Trëma* project was installed on the street at the entrance to the festival site, which required festivalgoers to pass through the installation to access the event. The producer also intended to use the artwork as a gateway to make a statement, given the scale and spectacular nature of the installation.

A pedestal was made available to the public, allowing interaction with the reinforcement learning algorithm that transforms the artwork's content. The minimalist wooden pedestal features a simple interface with three buttons. A green button rewards the system, signaling it to continue exploring within adjacent parametric combinations. The red button indicates that the participant does not particularly enjoy the current aesthetic proposal, prompting the agent to gradually shift toward other parametric combinations. A star button directs the agent to completely move to a different parametric zone.

The installation of the artwork consists of two large circular projection screens of 3.5 m in diameter, suspended on lighting truss structures. The circles are installed above the audience. We conceived the project not only as a generative audiovisual experience but, more broadly, as an expanded interactive urban scenography [13].

We used lasers to fill the space in all three dimensions. The laser imagery was generated from video patterns by tracing their contours. Fog machines revealed the beams of the video projectors, but more importantly, the laser rays, adding volumetric light [14] to the space. This created a cohesive scenographic object that visually connected the two screen structures. The sound was transmitted through four speakers placed around the installation area, significantly enhancing the immersive effect of the artwork.

Our aim was to transpose some of the controls we used in the laboratory to the general public, but the festive family context of the event proved unsuitable for this approach. Many participants, especially children, would frantically and randomly press buttons to observe the effects, without grasping the concept of interaction with the adaptive agent.

Despite this chaotic experience, the experiment proved rich in lessons. It revealed that the agent can offer some control in generating an immersive experience and highlighted the adjustments that could be made to ensure that the audience can fully engage with the work in future presentations. However, the way *Trëma* was presented in the public space introduced an element of unpredictable agency beyond our control. In particular, the interactive interface proved unsuitable for a festival setting where crowds of more than 5000 visitors per night wander around and through the installation. Reflecting on this experience, we concluded that the interface would be better suited for a more intimate context, such as an art gallery or museum, where the audience is more likely to engage deeply and individually with the work.

Regardless of the venue, it seemed necessary to rethink the role of audience members. Should they be seen as performers or DJs orchestrating the experience, or rather as art directors shaping the creative vision? Furthermore, we thought that the interface itself would benefit from a more deliberate scenographic presentation, perhaps with exclusive access or a better placed pedestal to highlight its significance and encourage thoughtful interaction.

2.3 Showcase in an Art Gallery

A new iteration was developed and artists were invited to co-perform with the installation. Our goal was to observe the system's performative potential in a more controlled context. With this objective in mind, we modified the interface to allow for finer control from the participant and help them visualize the agent's behavior beyond the generated audiovisual output. The pedestal was replaced with a tablet featuring a tactile control interface [9] designed using TouchOSC software[3]. In addition to the elements from the previous version, we added sliders for each audiovisual parameter controlled by the agent (see Figure 3). We also added toggle buttons to switch control between human and agent for any parameter. These parameters were also integrated into the visual dimension of the work by overlaying a graphical representation of each parameter's evolution onto the projection. This addition aimed to provide the observing audience with

[3]https://hexler.net/touchosc

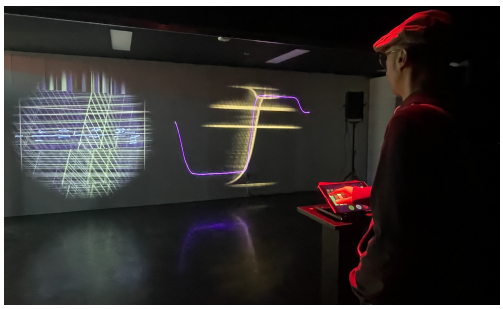

Figure 3: Second iteration of the project, presented in a gallery environment.

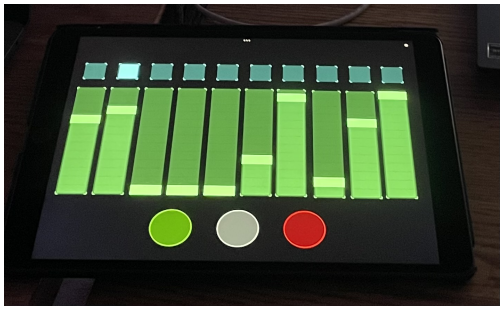

Figure 4: The second iteration of our interface.

an additional interpretative key to help understand the agent's behavior.

The installation setup was similar to the previous iteration, using projection onto two circles layered with lasers, a quadraphonic sound system and the tablet positioned on a stand facing it. Time constraints prevented us from inviting artists to perform. Instead, the audience engaged with the work, presented as an installation. This smaller audience consisted of researchers, artists, friends, and family members of the workshop participants.

The ability to view and take control of the audiovisual parameters changed the overall interaction behavior compared to previous iterations. The more explicit feedback reduced the appeal of the starred button used to radically switch parametric zones. We also observed that visualizing the parameters created more explicit expectations regarding the reinforcement learning behavior. These expectations stemmed from the participant's ability to configure the parameters in a specific way and then reward the learning process by repeatedly reinforcing that configuration. However, the limited interaction time and the multitude of contradictory interactions compromised the emergence of behaviors clearly adjusted to these reinforcements.

3 Discussion

3.1 From Performance to Installation

Traditional audiovisual performance practices are grounded in a model in which performers-creators carefully craft, sequence, and structure audiovisual media to produce an intentional narrative arc or emotional evolution over time. Creators act as both designers and performers, ensuring that each stage of the performance aligns with their artistic vision. The creative process revolves around composition, arrangement, and real-time execution, where the artist maintains direct and continuous control over the unfolding of the work.

The introduction of an IRL system like the Co-Explorer fundamentally disrupts this traditional paradigm. Instead of crafting predetermined content and transitions, the creative process becomes distributed across three interrelated components:

(1) Parametric Design: Human creators define the parametric space within which the RL agent operates. This involves selecting which audiovisual parameters the agent will control and determining the constraints that shape its behavior. (2) Agent Performance: The RL agent autonomously explores and generates content within the parametric space. Human creators observe and analyze the agent's behavior to understand its tendencies and characteristics. (3) Training: Human creators provide real-time positive or negative feedback, guiding the RL agent's evolution.

These three components not only redefine the creative process but also correspond to distinct roles that emerge from the fragmentation of traditional performer-creator roles. Rather than a single human individual or collective embodying both the designer and performer roles, these roles are now distributed among human and machine agents:

(1) Designer: Focused on constructing the parametric space, establishing the foundational structure within which the agent evolves to generate audiovisual content. (2) Performer: The RL agent itself takes the place of the performer, making autonomous decisions that are influenced by human co-creators. (3) Meta-performer: Acts as a selector, guide, or "gardener" [19], shaping the agent's behavior through subjective feedback without directly controlling the output.

Finally, these components do not remain static throughout the creative process, but instead evolves dynamically throughout the development of the artwork. The initial focus is largely on parametric design and observing the agent. Since the agent's behavior is still undeveloped, the creators focus on crafting an environment that offers both structure and flexibility. The agent's responses during this phase often reveal unforeseen possibilities, prompting adjustments to the design based on how the agent interacts with the parameters.

As the parametric space becomes more solid, attention shifts toward training the agent. The human role transitions from system designer to guide, reinforcing desirable behaviors and discouraging less interesting ones. This phase is marked by a growing relationship between the human and the agent, where the artist begins to intuitively understand the agent's tendencies and potential.

Eventually, the need for active training diminishes as the agent develops a coherent behavioral repertoire. At this point, the agent's autonomous performance becomes the focal point. The human creators step back, allowing the agent to operate independently within the framework, often uncovering emergent creative outcomes that were not explicitly programmed or anticipated.

This iterative evolution mirrors a process of co-adaptation, where both the human and the machine learn from each other. The agent's behavior influences how humans adjust the parametric space and provide feedback, while human interventions shape the agent's evolving performance capabilities.

3.2 Potential of Interactive Reinforcement Learning in the Development of Interactive Artworks

Interactive artworks often require immediate responsiveness for interpretability, yet RL unfolds over different timescales than

human decision-making. This creates a dichotomy: while interactivity demands instant feedback, RL functions through long-term adaptation. This tension raises questions about focus and immersion in duration-based works, as audiences must engage with evolving behaviors over extended periods.

The separation of performer roles in RL-based systems allows for new artistic configurations by distributing responsibilities to design, guide, and perform among different agents. This enables different configurations of RL-driven artworks leading to new forms of creative practices. Examples include:

- RL-Performer Album: A lineup of human **guides** is invited to each train a different RL agent based on a pre-**designed** parametric space: the resulting AI **performances** are played live by RL agents as a performance/installation/"album" hybrid.
- Interactive Parametric Space Showcase: A curated list of human **designers** is invited to each create a parametric space; the public is then invited to **guide** the RL-agents to generate different kinds of audiovisual experiences.
- Meta-DJ Night: A live performance where some audience members are invited to do a "meta-DJ" set where they get to act as the **meta-performer** while the rest of the audience is enjoying the show.
- Co-performance: A performance is **co-(meta)performed** live with human performers collaborating live with RL-agents, sometimes taking control and sometimes letting the agent take control.

These flexible assignments enable novel forms of co-creation between humans and machines, expanding the possibilities of interactive artistic experiences.

3.3 Strengths and Limitations of the Frameworks

The adaptability of reinforcement learning in artistic contexts remains constrained by technical requirements. The necessity for coding expertise creates a barrier to entry, limiting accessibility for artists who may not have a strong computer science background. Additionally, modifying or fine-tuning the system demands a deep understanding of both the algorithm and the specific artistic objectives, potentially restricting creative fluidity.

A challenge in using AI-driven tools in artistic practice is the risk of prioritizing the technology itself over artistic intent. In some cases, the integration of RL may feel externally imposed rather than organically embedded within the artist's creative process.

The way AI tools are designed inevitably shapes the resulting works. The parametric constraints set by the creators dictate the expressive possibilities, leading to a form of predetermined emergence rather than truly open-ended exploration. Drawing on Simondon's philosophy of individuation, we can consider the machine-learning system not merely as a tool but as a technical object with its own developmental trajectory [6]. The affordances and constraints of the chosen RL architecture play a defining role in the artistic process, underscoring the interdependence between technological mediation and aesthetic production.

4 Conclusion

It is pertinent to ask whether the integration of reinforcement learning really enhances the immersive experience or adds meaning to an artwork. Allowing the audience to interact with the system undeniably added value to the experience, especially if the aim is to reinforce the sense of immersion. The impression of exercising some control over the artwork seems to have an impact on the duration of participation.

However, interaction must be sufficiently significant from the outset to maintain interest and give the impression of active participation in the work. Yet, as we found in our initial experiments, developing a relationship with the artificial agent takes time. When participants pressed the positive feedback button, the agent registered it, but changes were not immediately visible. Sending repeated inputs slightly improved responsiveness, though effects remained subtle. The most impactful interaction was the "explore" button, which prompts the agent to move to a new area, producing immediate audiovisual effects. It was also necessary to regulate the use of this button, as participants tended to use it excessively to observe changes.

Our observation of the participants brought forward an important realization: the true value of reinforcement learning emerges over time. By the end of the evening, the agent had been shaped by several hours of interaction, revealing the work's full potential. Furthermore, its evolution across multiple evenings became an increasingly compelling aspect. However, because interactions occurred sporadically and often without sustained attention, the resulting behavior was also erratic. This highlighted the importance of considering the presentation context. The Lumifest digital arts festival, for instance, favors works suited to casual engagement, such as interactive projections that can be experienced briefly, with immediate and visible feedback. In contrast, the introduction of an artificial agent requires a context that supports more sustained and attentive interaction: one needs to spend time with the work to "get to know" it in an embodied and situated manner[1, p. 66]. In a gallery setting, for example, where the affordances of the work can signal the significance of participant involvement, the resulting behaviors tend to be more coherent and meaningful.

This research was primarily concerned with examining how the integration of machine learning transforms artistic practices, focusing on the creative process and the evolving role of the artist in dialogue with algorithmic agents. While audience interactions were briefly addressed (particularly in relation to the gallery exhibition), a more in-depth evaluation of public perception was beyond the scope of this study. We acknowledge that exploring how audiences interpret and respond to these co-created works, through qualitative interviews or surveys, would significantly enrich the analysis. This presents a valuable avenue for future research, particularly in understanding the broader cultural and aesthetic implications of machine learning in art.

5 Ethical Standards

This research was conducted in compliance with ethical guidelines. The authors and collaborators received compensation for their work, and the study was approved under an institutional ethics certificate. Additionally, we acknowledge and cite relevant open-source projects that contributed to the development of this work. The authors declare no conflicts of interest.

Acknowledgments

This work was supported by the Social Sciences and Humanities Research Council through the Insight Development program (funding #430-2022-00866), the Hexagram Network, and Hexagram-UQAM. Our laboratory, the mXlab, sits on ancestral and unceded Indigenous territory, a place known as Tiohtiá:ke

in Kanien'kéha and Mooniyang in Anishinaabemowin. We recognize that the Kanien'kehá:ka Nation are stewards of the lands and waters of this territory, which has long been a place where meetings and exchanges among nations have occurred.

References

[1] Sofian Audry. 2021. *Art in the Age of Learning Machines.* The MIT Press, Cambridge, Massachusetts.

[2] Margaret A. Boden. 2009. Computer Models of Creativity. *AI Magazine* 30, 3 (July 2009), 23–23. https://doi.org/10.1609/aimag.v30i3.2254

[3] Danielle Boutet. 2018. La Création de Soi Par Soi Dans La Recherche-Création : Comment La Réflexivité Augmente La Conscience et l'expérience de Soi. *Approches inductives* 5, 1 (2018), 289–310. https://doi.org/10.7202/1045161ar

[4] Oliver Bown. 2021. *Beyond the Creative Species: Making Machines That Make Art and Music.* The MIT Press, Cambridge, Massachusetts.

[5] Antoine Caillon and Philippe Esling. 2021. RAVE: A Variational Autoencoder for Fast and High-Quality Neural Audio Synthesis. https://doi.org/10.48550/ARXIV.2111.05011 arXiv:2111.05011 [cs]

[6] Pascal Chabot, Aliza Krefetz, and Graeme Kirkpatrick. 2013. *The Philosophy of Simondon: Between Technology and Individuation.* Bloomsbury Academic.

[7] Owen Chapman and Kim Sawchuk. 2015. Creation-as-Research: Critical Making in Complex Environments. *RACAR : Revue d'art canadienne* 40, 1 (2015), 49. https://doi.org/10.7202/1032753ar

[8] Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Douglas Eck, Karen Simonyan, and Mohammad Norouzi. 2017. Neural Audio Synthesis of Musical Notes with WaveNet Autoencoders. *arXiv:1704.01279 [cs]* (April 2017). arXiv:1704.01279 [cs]

[9] Günter Geiger. 2006. Using the Touch Screen as a Controller for Portable Computer Music Instruments. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME '06).* IRCAM — Centre Pompidou, 61–64.

[10] Michael Humphreys. 2005. Getting Personal: Reflexivity and Autoethnographic Vignettes. *Qualitative Inquiry* 11, 6 (Dec. 2005), 840–860. https://doi.org/10.1177/1077800404269425

[11] Vera John-Steiner. 2006. *Creative Collaboration* (pbk. ed ed.). Oxford University Press.

[12] Théo Jourdan and Baptiste Caramiaux. 2023. Machine Learning for Musical Expression: A Systematic Literature Review. 319–331 pages. https://doi.org/10.5281/zenodo.11189198

[13] Joslin McKinney and Scott Palmer (Eds.). 2017. *Scenography Expanded: An Introduction to Contemporary Performance Design.* Bloomsbury Publishing. https://doi.org/10.5040/9781474244428?locatt=label:secondary_dramaOnline

[14] Philippe-Alain Michaud. 2011. Line Light: The Geometric Cinema of Anthony McCall. 137 (2011), 3–22. https://doi.org/10.1162/OCTO_a_00055

[15] Néill O'Dwyer. 2021. *Digital Scenography: 30 Years of Experimentation and Innovation in Performance and Interactive Media.* Bloomsbury Methuen Drama.

[16] Louis-Claude Paquin and Cynthia Noury. 2020. Petit Récit de l'émergence de La Recherche-Création Médiatique à l'UQAM et Quelques Propositions Pour En Guider La Pratique. (2020), 103–136. Issue La communication à l'UQAM. https://doi.org/10.4000/communiquer.5042

[17] Hugo Scurto, Bavo Van Kerrebroeck, Baptiste Caramiaux, and Frédéric Bevilacqua. 2021. Designing Deep Reinforcement Learning for Human Parameter Exploration. *ACM Transactions on Computer-Human Interaction* 28, 1 (Feb. 2021), 1–35. https://doi.org/10.1145/3414472

[18] Richard S. Sutton and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction.* The MIT Press, Cambridge, Massachusetts.

[19] Stephen Todd and William Latham. 1992. *Evolutionary Art and Computers.* Academic Press, San Francisco, CA.

[20] Joanna Zylinska. 2020. *AI Art: Machine Visions and Warped Dreams.* Open Humanites Press.