

Designing With Machine Learning for Interactive Music Dispositifs

Hugo Scurto

▶ **To cite this version:**

Hugo Scurto. Designing With Machine Learning for Interactive Music Dispositifs. Sound [cs.SD]. Sorbonne Université, 2019. English. ⟨NNT : 2019SORUS356⟩. ⟨tel-02935929v2⟩

HAL Id: tel-02935929

https://hal.science/tel-02935929v2

Submitted on 12 Feb 2021

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.



HAL Authorization

Designing With Machine Learning for Interactive Music Dispositifs

Hugo SCURTO

THÈSE DE DOCTORAT DE
SORBONNE UNIVERSITÉ

École doctorale Informatique, Télécommunications et Électronique

Institut de Recherche et de Coordination Acoustique/Musique

Pour obtenir le grade de

DOCTEUR de SORBONNE UNIVERSITÉ

Spécialité **Informatique**

soutenue le 2 décembre 2019 devant le jury composé de :

Frédéric BEVILACQUA	Directeur de thèse	Ircam
Samuel BIANCHINI	Examinateur	École nationale supérieure des Arts Décoratifs
Sergi JORDÀ	Rapporteur	Universitat Pompeu Fabra
Wendy MACKAY	Rapportrice	Université Paris-Saclay
Catherine PÉLACHAUD	Examinatrice	Sorbonne Université
Anne SÈDES	Examinatrice	Université Paris 8

Hugo Scurto: *Designing With Machine Learning for Interactive Music Dispositifs*.
Final version as of January 2, 2020.

Supervisor: Frédéric Bevilacqua.

Abstract

Music is a cultural and creative practice that enables humans to express a variety of feelings and intentions through sound. Machine learning opens many prospects for designing human expression in interactive music systems. Yet, as a Computer Science discipline, machine learning remains mostly studied from an engineering sciences perspective, which often exclude humans and musical interaction from the loop of the created systems.

In this dissertation, I argue in favour of designing *with* machine learning for interactive music systems. I claim that machine learning must be first and foremost situated in human contexts to be researched and applied to the design of interactive music systems. I present four interdisciplinary studies that support this claim, using human-centred methods and model prototypes to design and apply machine learning to four situated musical tasks: motion-sound mapping, sonic exploration, synthesis exploration, and collective musical interaction.

Through these studies, I show that model prototyping helps envision designs of machine learning with human users before engaging in model engineering. I also show that the final human-centred machine learning systems not only helps humans create static musical artifacts, but supports dynamic processes of expression between humans and machines. I call co-expression these processes of musical interaction between humans—who may have an expressive and creative impetus regardless of their expertise—and machines—whose learning abilities may be perceived as expressive by humans.

In addition to these studies, I present five applications of the created model prototypes to the design of interactive music systems, which I publicly demonstrated in workshops, exhibitions, installations, and performances. Using a reflexive approach, I argue that the musical contributions enabled by such design practice with machine learning may ultimately complement the scientific contributions of human-centred machine learning. I claim that music research can thus be led through dispositif design, that is, through the technical realization of aesthetically-functioning artifacts that challenge cultural norms on music and computer science.

Résumé

La musique est une pratique culturelle permettant aux êtres humains d'exprimer sensiblement leurs intentions à travers le son. L'apprentissage machine définit un ensemble de modèles permettant de nouvelles formes d'expression au sein desdits systèmes interactifs musicaux. Cependant, en tant que discipline informatique, l'apprentissage machine demeure essentiellement appliquée à la musique du point de vue des sciences de l'ingénieur, qui, très souvent, conçoit les modèles d'apprentissage sans tenir compte des interactions musicales prenant place entre humains et systèmes.

Dans cette thèse, j'envisage la possibilité de mener des pratiques de design *avec* l'apprentissage machine pour les systèmes interactifs musicaux. Je soutiens que l'apprentissage machine doit avant tout être situé au sein d'un contexte humain afin d'être conçu et appliqué au design de systèmes interactifs musicaux. Pour défendre cette thèse, je présente quatre études interdisciplinaires, dans lesquelles j'introduis des modèles intermédiaires d'apprentissage, dits modèles-prototype, au sein de méthodes de conception centrées humain, afin d'appliquer l'apprentissage machine à quatre tâches musicales situées : le mapping mouvement-son, l'exploration sonore, l'exploration de la synthèse, et l'interaction musicale collective.

À travers ces études, je montre que les modèles-prototype permettent de générer des idées de design pour l'apprentissage machine en amont de la phase d'ingénierie desdits modèles, ce en lien étroit avec les utilisateurs potentiels de ces systèmes. Je montre également que les systèmes d'apprentissage machine centrés humain résultant de ce processus de conception rendent possible des processus dynamiques d'expression entre les humains et les machines, allant au-delà de la création d'artefacts musicaux statiques. Je propose de nommer co-expression ces processus d'interaction musicale entre des êtres humains—faisant preuve d'un élan expressif et créatif quelque soit leur expertise musicale—et des machines—dont les capacités d'apprentissage peuvent être perçues comme expressives du point de vue de l'humain.

En outre, je présente cinq systèmes interactifs musicaux conçus avec lesdits modèles-prototypes, et relate leurs restitutions publiques au sein d'ateliers, expositions, installations et performances. Par une approche réflexive, je montre que les contributions musicales apportées par des pratiques de design avec l'apprentissage machine peuvent, à terme, complémenter les contributions scientifiques apportées par les méthodes de conception centrées humain. Ainsi, je suggère que la recherche musicale peut être menée *par* le design de dispositifs interactifs musicaux, c'est-à-dire, par la réalisation technique d'artefacts esthétiquement fonctionnels remettant en cause les normes culturelles régissant l'informatique et la musique.

Acknowledgements

To my mentor Frédéric Bevilacqua, for genuinely tuning in to my creative impetus, and breathing harmonious lessons into the way—thank you for allowing me,

To artists and researchers I mixed with at Goldsmiths University of London, especially Rebecca Fiebrink, for being a source of inspiration for the time at work and beyond,

To scientists and physicists I befriended at École Normale Supérieure Paris-Saclay, especially Emmanuelle Deleporte, for helping thriving outside sinuous school years,

To musicians I played with at Cité de la Musique de Marseille, especially Patrick Vogel, for rocking out and acquainting with musical research and IRCAM in the first place,

To researchers and developers I studied with at IRCAM, especially in the ISMM group, for continually supporting and enriching thought and experimentation with sound,

To crafters and educators I engaged with at Lutherie Urbaine, especially Agnès Dufour, for rekindling a sense of materiality and a care for otherness in music,

To dancers and practitioners I moved with in the movA workshops, especially Anne Dubos, for ushering me in so many heightened scientific and artistic territories,

To countless friends and colleagues I worked with along this thesis, by any other means: Victor Audouze, Sofian Audry, Sylvie Benoit, Samuel Bianchini, Riccardo Borghesi, Baptiste Caramiaux, Djellal Chalabi, Axel Chemla–Romeu-Santos, Mélanie Egger, Philippe Esling, Sarah Fdili Alaoui, Emmanuel Flety, Jules Françoise, Jean-Louis Fréchin, Alain Guazzelli, Olivier Houix, Jean-Philippe Lambert, Joseph Larralde, Abby Wanyu Liu, Laurent Le Gall, Guillaume Lemaitre, Kevin Kamdoum, Michelle Agnes Magalhaes, Benjamin Matuszewski, Nicolas Misdariis, Uros Petrevski, Zelda Pollard, Norbert Schnell, Diemo Schwarz, Maël Segalen, Patrick Susini, Samuel Thompson, Bavo Van Kerrebroeck, Frédéric Voisin—and so many others by all means,

To my dearest musicians-designers at wolfgang, especially Axel Chemla–Romeu-Santos and Kévin Raharolahy, for always sustaining a child's ears,

To my people from Marseille and beyond, for ceaselessly stirring seas of sound,

To my parents, sister, and relatives, for bequeathing their forever love for music,

To Adèle Pécout, for being music herself.

Contents

Abstract iii

Résumé iv

Acknowledgements v

1 Introduction 1
- 1.1 Thesis Statement . . . 2
- 1.2 Research Approach . . . 3
  - 1.2.1 Situated Musical Tasks . . . 3
  - 1.2.2 Model Prototypes . . . 3
  - 1.2.3 Human-Centred Methods . . . 4
  - 1.2.4 Reflexive Approach . . . 5
- 1.3 Contributions . . . 5
  - 1.3.1 Model Prototypes . . . 5
  - 1.3.2 Empirical Findings . . . 6
  - 1.3.3 Music Dispositifs . . . 6
  - 1.3.4 Theoretical Perspectives . . . 7
- 1.4 Thesis Overview . . . 7

2 Context and Related Work 11
- 2.1 Interactive Music Systems . . . 11
  - 2.1.1 On Music Practices . . . 11
  - 2.1.2 Design Approaches . . . 13
  - 2.1.3 Interaction Techniques . . . 14
- 2.2 Machine Learning . . . 16
  - 2.2.1 Formal Definitions . . . 17
  - 2.2.2 Musical Applications . . . 18
  - 2.2.3 On Research Methods . . . 20
- 2.3 Human-Centred Machine Learning . . . 22
  - 2.3.1 Motivation . . . 22
  - 2.3.2 Research Approach . . . 23
  - 2.3.3 Example Musical Applications . . . 24
- 2.4 Summary . . . 26

3 Designing With Unsupervised Learning For Motion-Sound Mapping 27
3.1 Concepts . . . 28
3.1.1 Musical Task: Motion-Sound Mapping . . . 28
3.1.2 Proposed Technique: Unsupervised Learning . . . 30
3.2 Prototyping . . . 30
3.2.1 Algorithm: *Online Gaussian Mixture Model* . . . 30
3.2.2 Human Data: Observation . . . 31
3.2.3 Implementation: Patch . . . 33
3.3 Application: *somasticks* . . . 34
3.3.1 Description . . . 34
3.3.2 Instrument Design . . . 34
3.3.3 Demonstration in Workshop . . . 37
3.4 Discussion . . . 39
3.4.1 Challenges in Implementation . . . 39
3.4.2 Designing for Human Exploration . . . 40

4 Designing With Reinforcement Learning For Sonic Exploration 41
4.1 Concepts . . . 42
4.1.1 Musical Task: Sonic Exploration . . . 44
4.1.2 Proposed Technique: Reinforcement Learning . . . 46
4.2 Prototyping . . . 48
4.2.1 Human Data: Feedback . . . 48
4.2.2 Algorithm: *Sarsa* . . . 49
4.3 Observation . . . 49
4.3.1 Setup: Controlled Experiment . . . 49
4.3.2 Results: Statistical Analysis . . . 52
4.4 Discussion . . . 55
4.4.1 The Usefulness of Balancing Exploitation with Exploration . . . 55
4.4.2 The Influence of Exploration Path on User Perception . . . 56
4.4.3 Towards Co-Exploration . . . 57

5 Designing With Deep Reinforcement Learning For Synthesis Exploration 59
5.1 Concepts . . . 60
5.1.1 Musical Task: Synthesis Exploration . . . 61
5.1.2 Proposed Technique: Deep Reinforcement Learning . . . 64
5.2 Prototyping I . . . 65
5.2.1 Algorithm: *Sarsa* . . . 65
5.2.2 Human Data: Feedback . . . 67
5.3 Observation I . . . 67
5.3.1 Setup: Pilot Study . . . 67
5.3.2 Results: VST Interfaces, Qualitative Analysis . . . 67
5.3.3 Results: RL Agent Prototype, Qualitative Analysis . . . 70
5.4 Prototyping II . . . 72
5.4.1 Human Data: Feedback + Controls . . . 72
5.4.2 Algorithm: *Deep TAMER* . . . 74

5.5 Engineering II . . . 75
5.5.1 Algorithm: Improving Parameter Exploration . . . 75
5.5.2 Human Data: Processing Feedback and Controls . . . 75
5.5.3 Implementation: *Co-Explorer* . . . 77
5.6 Observation II . . . 78
5.6.1 Setup: Workshop . . . 78
5.6.2 Results: Discovery Task, Qualitative Analysis . . . 78
5.6.3 Results: Creation Task, Qualitative Analysis . . . 81
5.7 Discussion . . . 83
5.7.1 Conceptual Insight . . . 84
5.7.2 Technical Insight . . . 85
5.7.3 Guidelines for Designing Machine Learning for Creativity . . . 86
5.8 Application I: *Riding the Co-Explorers* . . . 89
5.8.1 Description . . . 89
5.8.2 Game Design . . . 90
5.8.3 Application in Exhibition . . . 91
5.9 Application II: *Behavioral Matter* . . . 93
5.9.1 Description . . . 93
5.9.2 Object Design . . . 93
5.9.3 Application in Workshop . . . 95
5.10 Application III: *ægo* . . . 96
5.10.1 Motivation . . . 96
5.10.2 Scientific Modelling . . . 97
5.10.3 Instrument Design . . . 99
5.10.4 Musical Artwork . . . 101
5.10.5 Discussion . . . 103

6 Designing With Active Learning For Collective Musical Interaction 107
6.1 Observation . . . 108
6.1.1 Musical Task: Collective Musical Interaction . . . 108
6.1.2 Setup: Field Study . . . 109
6.1.3 Results: Qualitative Analysis . . . 110
6.2 Prototyping . . . 112
6.2.1 Proposed Technique: Active Learning . . . 112
6.2.2 Human Data: Performance . . . 112
6.2.3 Algorithm: *Bayesian Information Gain* . . . 113
6.2.4 Implementation: Web-based Application . . . 114
6.3 Application: *entrain* . . . 114
6.3.1 Description . . . 114
6.3.2 Installation Design . . . 115
6.3.3 Application in Exhibition . . . 116

7 On Human-Centred Machine Learning Applied to Music **121**
7.1 Harnessing Human Musical Goals and Values . . . 121
7.1.1 Human Feedback on Musical Tasks . . . 121
7.1.2 Music Practices and Machine Learning . . . 123
7.2 Switching Machine Learning Tasks for Interaction Techniques . . . 125
7.2.1 Model Prototypes as Design Artifacts . . . 125
7.2.2 Reassigning Model Engineering In The Design Process . . . 127
7.3 Creating Co-Expressive Human-Machine Practices . . . 129
7.3.1 Machine Expression . . . 130
7.3.2 Human-Machine Co-Expression . . . 132
7.3.3 On Human Expression, Music, and Artificial Creativity . . . 135

8 Conclusion **137**
8.1 Summary and Contributions . . . 137
8.2 Directions for Future Work . . . 138
8.2.1 Human-Centred Perspectives . . . 139
8.2.2 Machine-Centred Perspectives . . . 139
8.2.3 Musical Perspectives . . . 140

Post-Scriptum: Music Research Through Dispositif Design **143**
Doctoral Research in Science Applied to Music . . . 143
Music Research Through Design . . . 144
From Interactive Music Systems to *Interactive Music Dispositifs* . . . 146
Norms in Music . . . 146
Norms in Machine Learning . . . 148
Example Projects . . . 150
Sound Control . . . 151
Paroles d'Exil . . . 152
{Lutheries}² . . . 154
Music (or) Design Practice? . . . 156

Publications **161**

Videography **163**

Bibliography **165**

List of Figures

Figure 1.1 The four situated musical tasks investigated in this thesis. . . 3
Figure 1.2 The four model prototypes designed in this thesis. 4
Figure 1.3 The four human-centred methods used in this thesis. 4
Figure 1.4 The real-world applications of our model prototypes. 5
Figure 1.5 Graphical outline of the four studies led in the thesis. 9

Figure 3.1 The design-oriented method for our first study. 27
Figure 3.2 Interactive workflow for *Mapping-by-Demonstration*. 29
Figure 3.3 Graphical outline of a *Gaussian Mixture Model*. 31
Figure 3.4 Interactive workflow for *Design Through Performance*. 32
Figure 3.5 Graphical outline of our two interaction modes. 32
Figure 3.6 User interface of our model prototype. 34
Figure 3.7 Rendering for the 3D-printed support for *somasticks*. 35
Figure 3.8 Graphical outline of the mapping used in *somasticks*. 36
Figure 3.9 Picture of a performer playing with *somasticks*. 37

Figure 4.1 The experimental method for our second study. 41
Figure 4.2 Example user interface of a music computing system. 43
Figure 4.3 Graphical outline of *Interactive Reinforcement Learning*. 47
Figure 4.4 Interactive workflow for *Co-Exploration*. 48
Figure 4.5 Graphical outline of the two sonic environments of our study. 51
Figure 4.6 Results for synthetic trial data. 52
Figure 4.7 Results for participants' trial data. 53
Figure 4.8 Results for participants' evaluation data. 54
Figure 4.9 Results for participants' ratings versus task parameters. 55

Figure 5.1 The user-centered design method for our third study. 59
Figure 5.2 Example user interface of a VST for sound design. 62
Figure 5.3 Graphical outline of our RL agent prototype. 66
Figure 5.4 Graphical outline of the three parametric interfaces of our study. 68
Figure 5.5 Results for human synthesis exploration strategies. 69
Figure 5.6 Picture of a sound designer using our RL agent prototype. . . 70
Figure 5.7 Interactive workflow of the *Co-Explorer*. 73
Figure 5.8 Graphical outline of machine exploration methods. 76
Figure 5.9 Graphical outline of human feedback computation methods. . 76

Figure 5.10 User interface of the *Co-Explorer*. 77
Figure 5.11 Picture of sound designers using the *Co-Explorer* in our workshop. 78
Figure 5.12 Picture of the four images used in our workshop. 81
Figure 5.13 Results for human-machine co-exploration partnerships. . . . 82
Figure 5.14 Visual theme for *Riding the Co-Explorers*. 89
Figure 5.15 Screenshot taken from *Riding the Co-Explorers*. 90
Figure 5.16 User interface for *Riding the Co-Explorers*. 91
Figure 5.17 Picture of the exhibition stand used at *ISMIR'18*. 92
Figure 5.18 Picture of the *Behavioral Matter* workshop at *Centre Pompidou*. 94
Figure 5.19 Picture of the robotic object prototyped during the workshop. 95
Figure 5.20 Graphical outline of a *Variational Auto-Encoder*. 97
Figure 5.21 Graphical outline of *Deep Reinforcement Learning*. 98
Figure 5.22 Interactive workflow for our musical instrument. 100
Figure 5.23 Implementation of our musical instrument. 100
Figure 5.24 Pictures taken from the *ægo* artwork. 102
Figure 5.25 Graphical outline of *ægo* temporal structure. 102
Figure 5.26 Graphical outline of our research and creation process. 104

Figure 6.1 The participatory design method for our fourth study. 107
Figure 6.2 User interfaces for the *Coloop* application. 109
Figure 6.3 Graphical outline of *Bayesian Information Gain*. 113
Figure 6.4 Picture of the *Coloop mini* loudspeaker designed by *Nodesign.net*. 115
Figure 6.5 Picture of the *entrain* installation at *SIGGRAPH'19*. 117

Figure 7.1 Graphical outline of machine expression (situated in machines). 130
Figure 7.2 Graphical outline of machine expression (perceived by humans). 132
Figure 7.3 Graphical outline of co-expression (with machine learning). . 134

Figure A Graphical outline of normative elements in music. 147
Figure B Graphical outline of normative elements in machine learning. 149
Figure C Picture of the *Sound Control* project at *Northampton, UK*. . . . 151
Figure D User interface for the *Sound Control* dispositif. 152
Figure E Pictures taken from *Paroles d'Exil* at *Friche Belle de Mai*. 153
Figure F Control interface for the *Mémoire* dispositif. 154
Figure G Pictures of the *{Lutheries}*^{2} exhibition at *Bagnolet, France*. . . . 155
Figure H Picture of the music dispositif designed for the *Digitalities* area. 155
Figure I Graphical outline of the shifting of norms through design. . . 158

Glossary

BIG Bayesian Information Gain.

GMM Gaussian Mixture Model.

HCI Human-Computer Interaction.

ML Machine Learning.

NIME New Interfaces for Musical Expression.

RL Reinforcement Learning.

Sarsa State-Action-Reward-State-Action.

TAMER Training an Agent Manually via Evaluative Reinforcement.

VAE Variational Auto-Encoder.

VST Virtual Studio Technology.

alt sem hann sér
alt sem hún do

— Björk.

1

Introduction

Music is a cultural and creative *practice* that enables humans to express a variety of feelings and intentions, especially by means of embodied interaction with sound [Leman, 2016]. Different types of music instruments were created by humans in different cultural contexts to support different musical expression [Jorda, 2005]. Acoustic instruments have been used for centuries over the world to support the transformation of human gestures into sound. Electronic instruments were invented over the last century in Western countries, enabling new forms of sound generation by means of electronic and digital components. Specifically, *interactive music systems* were created over the last half century by computer music research communities to investigate novel ways to control sound, possibly blending digital software with physical components. We are interested in designing interactive music systems that takes into account existing music practices to enable new forms of human expression.

Machine learning opens many prospects for the design of interactive music systems. The framework supports the computational performance of many different tasks related to music practice, such as gestural control of sound [Bevilacqua et al., 2009], sound generation [Esling et al., 2019], or human-machine improvization [Assayag et al., 2006]. Machine learning is powerful because it automatically learns to perform the above-mentioned tasks from example data. For example, imagine that you could build a custom gestural controller by only recording examples of gestures and without doing any programming. The so-called mapping by demonstration [Françoise, 2015] is one example of expressive interaction enabled by machine learning in interactive music systems.

As a Computer Science discipline, machine learning remains mostly designed and applied to music from an *engineering sciences* perspective. Much effort is put into the solving of technical problems related to the non-interactive learning of musical tasks, such as automatic music analysis and generation. This approach may often exclude human users from the loop of the created machine learning system [Briot et al., 2017]. While engineering sciences methods remain essential to the building of effective machine learning systems, we believe that fully-automated systems may not always benefit humans in their music practices. Human-centred methods may be required to balance model engineering with human evaluation to truly foster human

creativity over artificial creativity [Beaudouin-Lafon and Mackay, 2018].

Human-centred machine learning is an interdisciplinary area of work that aims at rethinking machine learning research and design in terms of human goals [Gillies et al., 2016, Ramos et al., 2019]. The area borrows many methods from the field of Human-Computer Interaction to design machine learning systems that include humans in their loop. For example, rather than fully autonomous, these systems may let humans parametrize algorithms, or interactively provide example data to adapt machine learning to their needs. Previous works showed the usefulness of human-centred machine learning for the design of interactive music systems [Fiebrink, 2011, Françoise, 2015, Scurto, 2016]. Specifically, musicians may twist the initially-designed machine learning task to fulfill their musical tasks in a customized and creative way.

We are interested in *applying human-centred machine learning to the design of interactive music systems*. We expect that this methodology would enable to build a scientific understanding of the technical requirements of machine learning in relation to creative practices of music. Alternatively, we also envision to adopt a reverse approach, which is to *practice with machine learning for real-world musical situations*. We hypothesize that such a creative approach would enable to investigate other inquiries related to machine learning applied to music, in a way similar to how *dispositifs* questions cutting-edge technologies in the fields of art and design [Bianchini and Verhagen, 2016]. Overall, we expect that our joint scientific and practical approach will help us get a holistic understanding of machine learning in relation to music [Jorda, 2005].

1.1 Thesis Statement

In this dissertation, I argue in favour of designing *with* machine learning for interactive music systems. I claim that machine learning must be first and foremost situated in human contexts to be researched and applied to the design of interactive music systems. I present four interdisciplinary studies that support this claim, using *human-centred methods* and *model prototypes* to design and apply machine learning to four *situated musical tasks*: motion-sound mapping, sonic exploration, synthesis exploration, and collective musical interaction. Through these studies, I show that model prototypes help envision designs of machine learning with human users before engaging in engineering. I also show that the final human-centred machine learning systems not only helps humans create static musical artifacts, but supports dynamic processes of expression between humans and machines. I call *co-expression* these processes of musical interaction between humans—who may have an expressive and creative impetus regardless of their expertise—and machines—whose learning abilities may be perceived as expressive by humans. Using a *reflexive approach* based on real-world applications of the models, I argue that the musical contributions enabled by design practice with machine learning may ultimately complement the scientific contributions of human-centred machine learning. I claim that *music research* can thus be led *through dispositif design*, that is, through the technical realization of aesthetically-functioning artifacts that challenge cultural norms on music and computer science.

1.2 Research Approach

The research approach to human-centred machine learning applied to music consisted in the following loop: (1) focus on one situated music task, (2) rely on model prototypes to study one machine learning technique, and (3) use one human-centred method to research and design human interaction with machine learning. We iterated four times over this loop, each time leveraging a new musical task, machine learning technique, and human-centred method. In parallel, I also practiced with machine learning to create interactive dispositifs in real-world musical situations. This process enabled me to (4) adopt a reflexive approach to human-centred machine learning applied to music. In a post-scriptum added at the end of this dissertation, I relate my personal reflections on applied and practice-based approaches to machine learning for music, discussing disciplinary frameworks embedded in machine learning as well as in music.

1.2.1 Situated Musical Tasks

We applied human-centred machine learning to *situated musical tasks*. By situated, we mean taking into account the cultural, individual, or industrial context of music, as an embodied act of human expression [Leman, 2016]. Importantly, a given musical task may pertain to several music communities or practices. For example, the task of motion-sound mapping mostly pertains to performers from the NIME community. Alternatively, the task of sonic exploration may pertain to sound designers and composers from different music practices [Garcia et al., 2012], but also to non-musicians—*e.g.,* novice users learning to use a new interface [Resnick, 2007]. Situating musical tasks would not only enable us to take into account differences in music expertise, but also cultural differences in music practices—supporting the human-centred design of machine learning [Gillies et al., 2016]. We thus led our research and design process on four situated musical tasks: motion-sound mapping, sonic exploration, synthesis exploration, and collective musical interaction (see Fig. 1.1).

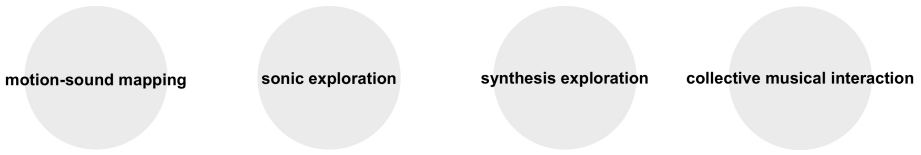


Figure 1.1: The four situated musical tasks investigated in this thesis.

1.2.2 Model Prototypes

We relied on *model prototypes* to study machine learning techniques in relation with the situated musical tasks. Model prototypes may be considered as design artifacts for machine learning techniques in interactive systems. They enable to test interactive data workflows with concrete algorithmic implementations before starting the engineering of a final learning model. As such, model prototypes enable to generate new

ideas and envision new designs of machine learning in situation with human users. This contrasts with engineering sciences approaches to machine learning, which often consider user interaction only after a model is optimized from a large data set.

The notion of model prototype extend that of software prototype [Beaudouin-Lafon and Mackay, 2009] to the case of statistical models in machine learning. For example, one may test several model prototypes—*e.g.*, centroid- or density-based—to design the machine learning technique of clustering [Murphy, 2012]. Focusing on model prototyping over model engineering enabled us to study four machine learning techniques in relation to the four situated musical tasks: unsupervised learning, reinforcement learning, deep reinforcement learning, and active learning (see Fig. 1.2).

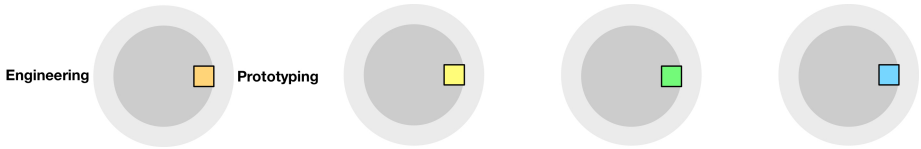


Figure 1.2: The four model prototypes of the thesis, respectively designed with unsupervised, reinforcement, deep reinforcement, and active learning.

1.2.3 Human-Centred Methods

We used *human-centred methods* to research and design our model prototypes in situated musical tasks. Human-centred methods come from the interdisciplinary field of Human-Computer Interaction (HCI), which builds interactive systems to observe human interaction and contribute to a theoretical knowledge of human-computer interaction [Mackay and Fayard, 1997]. Instead of theory, our human-centred methods will build on the interaction design research notion of *strong concepts*—an intermediate form of knowledge that is more abstracted than particular instances, yet does not aim at building general theories [Höök and Löwgren, 2012]. As such, we relied on four human-centred methods to iterate between **concepts**, **observation**, **prototyping**, and **engineering** with our model prototypes: design-oriented, experimental, user-centered, and participatory design methods (see Fig. 1.3).

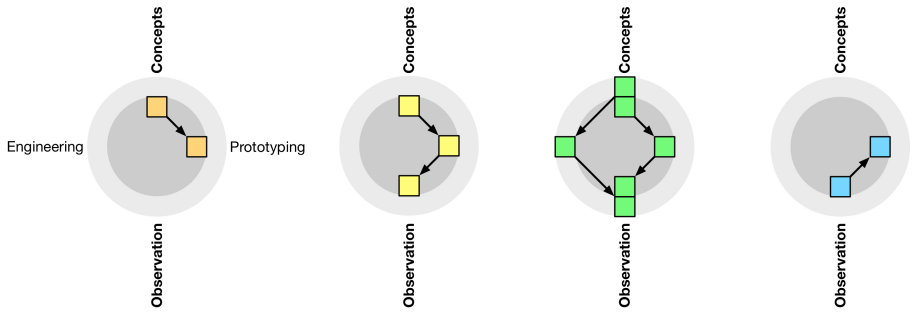


Figure 1.3: The four human-centred methods used in the thesis: design-oriented, experimental, user-centered, and participatory design methods.

1.2.4 Reflexive Approach

We adopted a *reflexive approach* to the application of human-centred machine learning to interactive music systems [Kläy et al., 2015]. In addition to concepts, design, observation, and engineering of human-centred methods, we led **applications** of our model prototypes in real-world musical situations (see Fig. 1.4). While human-centred methods have the scientific goal of evaluating the efficiency of a design, applications in real-world musical situations have the creative goal of sharing the experience and the conceptual ideas of a system with people and society in general. In this sense, they helped us adopt a reflexive approach by taking a critical stance on human-centred methods and explore practice-based approaches to music research.

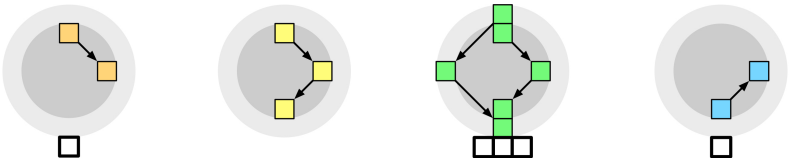


Figure 1.4: The real-world applications of our model prototypes, led outside the frame of the human-centred methods.

1.3 Contributions

The contributions of the thesis situate at four levels: (1) model prototypes, (2) empirical findings, (3) music dispositifs, and (4) theoretical perspectives.

1.3.1 Model Prototypes

Each of our four human-centred studies led to the design of an interactive model prototype, applying a machine learning technique to a situated musical task:

- The *Online Gaussian Mixture Model* model supports the musical task of motion-sound mapping by implementing unsupervised learning from human motion observations.
- The *Sarsa* model supports the musical task of sonic exploration by implementing reinforcement learning from human feedback and machine exploration.
- The *Deep TAMER* model supports the musical task of synthesis exploration by implementing deep reinforcement learning from human feedback, control and machine intrinsically-motivated exploration.
- The *Bayesian Information Gain* model supports the musical task of collective musical interaction by implementing active learning from human individual performance.

1.3.2 Empirical Findings

Observing humans interact with our four model prototypes enabled us to harvest empirical findings related to the musical tasks and situated in music practices:

- Performers projected themselves into imaginary musical situations interacting with unsupervised learning. This projection was enabled by the online adaptative behaviour of the motion-sound mapping.
- Humans perceive differences between various exploration strategies of reinforcement learning agents. These differences depend on environmental factors related to timbral features of sound.
- Sound designers partner with deep reinforcement learning agents in synthesis exploration. These partnerships can be placed along a continuum ranging from user- to agent-as-leader.
- Non-musicians reacted positively to the collective musical behaviour of active learning. This reaction was enabled by the expressive audiovisual queries made by the machine to the humans.

1.3.3 Music Dispositifs

In addition to our model prototypes and observations, we created several music dispositifs throughout out work. We propose to borrow the term "dispositif" from the fields of contemporary art and design [Bianchini and Verhagen, 2016] to consider broader dimensions of interactive music systems—we will properly introduce it in the Post-Scriptum. We applied the interactive dispositifs in several real-world musical situations, ranging from public exhibitions, performances, installations, and workshops.

- The *somasticks* are augmented drumsticks that use unsupervised learning to emphasize somatic expression in drumming practice.
- The *Co-Explorer* is a software that lets humans explore large parameter spaces by expressing positive or negative feedback to a deep reinforcement learning agent.
  - *Riding the Co-Explorers* is a single-level game that lets a single player guide the Co-Explorer in one single sound space.
  - *Behavioral Matter* is an interdisciplinary workshop that applied the Co-Explorer to the design of robotic objects.
  - *ægo* is an improvisational computer music piece with interactive sound and image for one performer and the Co-Explorer.
- *entrain* is a public installation that lets humans express collectively in music in collaboration with an active learning loudspeaker.

1.3.4 Theoretical Perspectives

Our human-centred studies and real-world applications enabled us to open theoretical perspectives on music, machine learning, research, and design.

- The artifacts called *model prototypes* help envision possible designs of machine learning in situation with human users before fully engineering them.
- The notions of *machine expression* and *co-expression* support the general analysis of musical interaction between (possibly learning) machines and (possibly non-musician) humans.
- The concept of *interactive music dispositif* supports explicit consideration of the scientific and musical norms that may condition human imagination on interactive systems.
- The approach of *music research through design* enables practice-based contributions in music that may complement applicative contributions in computer science by challenging the norms in interactive music systems.

1.4 Thesis Overview

Chapter 2 relates the context and related work for the thesis. It reviews basic notions on interactive music systems, machine learning, and human-centred machine learning applied to music.

Chapter 3 presents the first study of the thesis. It uses a design-oriented method to apply unsupervised learning to the musical task of motion-sound mapping. It describes the musical and technical concepts surrounding the study, followed by the design of a model prototype, and the application to the design of a music dispositif, called *somasticks*, that was tested by expert performers.

Chapter 4 presents the second study of the thesis. It uses an experimental method to apply reinforcement learning to the musical task of sonic exploration. It describes the musical and technical concepts surrounding the study, as well as the design of a model prototype, and the observation of a group of humans, containing both musicians and non-musicians, interacting with the model in a controlled setup.

Chapter 5 presents the third study of the thesis. It uses a user-centered design method to apply unsupervised learning to the musical task of synthesis exploration. It describes the musical and technical concepts surrounding the study, followed by the iterative design of a model prototype, intertwined with observations with expert sound designers, and ended by the engineering of the model. Lastly, it reports on applications of the model to the creation of three music dispositifs for real-world situations: *Riding the Co-Explorers*, *Behavioral Matter*, and *ægo*.

Chapter 6 presents the fourth study of the thesis. It uses a participatory design method to apply active learning to the musical task of collective musical interaction. It describes observations made with designers of previous collaborative mobile music-making systems, followed by the design of a model prototype, and the application to the design of a music dispositif, called *entrain*, that was tested by non-musicians.

Chapter 7 discusses the human-centred approach to machine learning applied to music adopted in the thesis. It describes the human and machine advantages of situating machine learning in human contexts to study musical interaction. It introduces the notions of *machine expression* and *co-expression* to describe such diversity of musical interaction between humans and learning machines. It builds on the latter human-centred notions to discuss artificial creativity approaches to machine learning applied to music.

The Post-Scriptum builds on real-world musical situations of practice with machine learning to reflect on the research led in the thesis. It discusses the balances made between intersecting perspectives from Computer Science and music. It presents design projects, led in parallel of the thesis, that challenge the norms of these disciplines. It finally delineates a *music research through design* framework, which leverage the theoretical notion of *interactive music dispositif* to foster novel musical practices that are grounded in their environment.

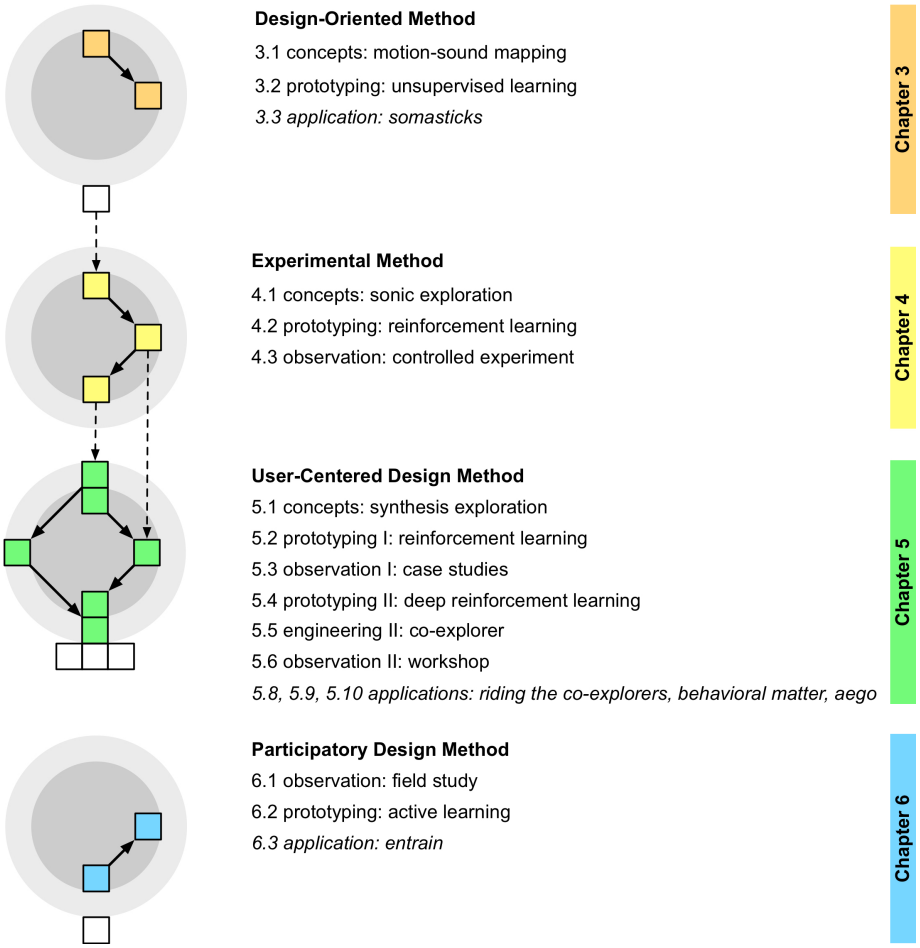


Figure 1.5: Graphical outline of the four studies led in the thesis.

2

Context and Related Work

In this chapter, we motivate our research approach by situating it in the context of related works. We start by contextualizing interactive music systems, situating the musical values that they may encapsulate, and describing the approaches and techniques used to design them. We then give an overview of machine learning, defining the technical elements of the framework, and reviewing applications and methods adopted in the field of music. We finally motivate human-centred machine learning, as a research approach that explicitly consider humans in the design of machine learning. We describe the techniques and methods employed, notably for applications to the field of music.

Part of this work was published and presented as doctoral symposium paper at the *4th International Conference on Movement Computing (MOCO 2017)* in London, United Kingdom [Scurto and Bevilacqua, 2017].

2.1 Interactive Music Systems

Interactive music systems are a broad family of digital systems aimed at supporting human music practices. These practices may encompass music performance, composition, and improvization, as well as listening, pedagogy, or playing. This section aims at describing the context in which the research and development of interactive music systems may operate. We first situate interactive music systems among a variety of musical values, which cannot be dismissed when considering human musical practices. We then describe the approaches adopted to their design, and the technical features that have been shown to be central to their functioning.

2.1.1 On Music Practices

As support of human musical practices, interactive music systems always encapsulate musical values and signification into their design. In this section, we situate our approach among a set of musical values. We argue in favour of (1) practical and (2) embodied dimensions of music, which we do by reviewing related work in musicology and cognitive science.

Music Is A Practice, Not A Thing

To borrow the expression of Christopher Small, music may be thought of as a *practice*, not a thing. According to him, there is no such "thing" called "music", but an abstraction of the human "action whose reality vanishes as soon as we examine it too closely" [Small, 1998]. He introduced the verb "musicking" to highlight the very fact that music may only exist through practice. For several decades, the signification of music was mostly assessed by relying on Western music theory—*e.g.,* the formalization of tones, harmony, and rhythm in an annotated form. The application of one single music theory over the variety of music practices tended to fix music as a thing, as an object that can be talked from an external point of view, without participation of its subject. Many musicologists also argued in favour of the so-called "performative turn" in music [Jost, 2013], whether it be acted through listening [Donin, 2004], or performing [Leman, 2008]. For example, Simon Frith argues that the signification and value of music may arise from rituals of identification made during listening by humans—perhars more than from the formal constructs of music [Frith, 1998].

Over the centuries, new forms of music practices were enabled by the introduction of technology. For example, phonograph records enabled humans to listen to music at a different place than where it used to be played—as is extended with latest streaming services. Slit drums enabled humans to perform music by creating complex rhythms in interaction with other humans—as is extended with latest smartphone applications. Because of their technological roots, many of these music practices may thus be linked with the industry [Wikström, 2013]. Although rejected by contemporary music institutions decades ago [Born, 1995], the links between public and private industrial sectors of music are currently subject to important transformations. Unfortunately, debating political issues on the industrialization of music remains out of the scope of this thesis. Yet, our wish was to stress that industry links must be considered when situating music practice in a societal context.

Music Is Rooted In Embodiment and Expression

The performative aspect of music may be apprehended from a cognitive point of view through the notion of *embodiment*. In his seminal book on embodied music cognition, Marc Leman argues that the human body may act as the main mediator between musical experience (mind) and sound energy (matter) [Leman, 2008]. Practices such as listening to music may be understood as an action-oriented phenomenon, which largely relies on our sensori-motor systems. For example, humans may move along the music they are listening to, to help make sense of auditory information [Nymoen et al., 2011]. Reciprocally, when playing a music instrument, humans may use their body to encode an idea, or mental representation, into a material or energetic form—*i.e.,* sound [Godøy and Leman, 2010]. The so-called musical gestures may be analysed and classified in a gesture space depending on their role in music performance—*i.e.,* sound-producing, sound-accompanying, or communicative gestures [Cadoz and Wanderley, 2000].

Embodied interaction with music is deeply linked with the notion of *human expression* [Leman, 2016]. The encoding of gestures into sound may be analysed as the

metaphorical expression of some human intention [Caramiaux et al., 2014a, Lemaitre et al., 2017]. Reciprocally, the decoding of sounds into gestures may help interpret the expression of a musical intention—*e.g.,* a crescendo being perceived as the acceleration of a human movement [Leman, 2016]. As such, encouraging musical expression may be crucial to support the social and emotional development of human beings, be they musicians, or non-musicians. Notions are as diverse as empathy [Altenmüller et al., 2013], attention [Hannon and Trainor, 2007], adaptation [Schlaug et al., 2005], self-esteem [Henderson, 1983], or creativity [Sawyer, 2014].

Music is an embodied practice that needs situating in a human context. This context may include industrial, cultural, and individual dimensions. Our work will build on culturally- and individually-situated music practices to design interactive music systems for human expression.

2.1.2 Design Approaches

In this section, we describe the design approaches that may be adopted to the building of interactive music systems. We present the corresponding (1) engineering sciences and (2) crafting approaches by reviewing the related work in computer music.

The Engineering Sciences of Interactive Music Systems

Engineering sciences support the fine-tuned design of individual components of interactive music systems. Focusing on individual components can enable bottom-up technique discoveries that may support new music practices. For example, many of the first synthesizers were created by electrical engineers from manipulation of analog electromechanical circuits [Bode, 1984]. Also, it can be used as a way to address the complexity of both digital and physical components in interactive music systems. For example, the engineering sciences background of Pierre Schaeffer arguably contributed to the development of his practice of sound and music [Schaeffer, 1966]. In addition, engineering sciences approaches may go in line with industrial applications that require system robustness for large-scale music production.

Yet, only relying on engineering sciences can cause harm to music practices. Taking technical components outside the scope of human interaction may direct the design of interactive music systems toward machines. While the design of music machines can be an exciting intellectual avenue [Rowe, 1992], it has been described as potentially enclosing fields of knowledge related to music [Sterne, 2003]. For example, the engineering of interactive systems based on implicit notions of "musical input" may restrict music performance toward practices based on Western music theory [Jorda, 2005]. The introduction of the computer may potentially increase this gap, by relying on digital computation over human actions. For example, the engineering of music recommender systems may direct music listening toward a practice of automated consumption over that of self-identification [Taylor, 2014]. All in all, the externalization of the machine from the human may contribute to the thing-ification of music over human music practices.

The Crafting of Computers and Materials

On the other hand, crafting supports the holistic design of digital and physical components into interactive music systems [Jorda, 2005]. It makes it suitable to some form of top-down design of interactive music systems, where the discovery of new techniques may be enabled by intuition emerging from musical practices. For example, Maurice Martenot, the inventor of the ondes Martenot, and also a cellist, was inspired by the accidental overlaps of tones between military radio oscillators, and wanted to create an instrument with the expressiveness of the cello [Leipp, 1972]. The practical aspect of crafting may go beyond the physical to embrace computers and digital components. Michel Waiswisz' The Hands may be the most well-known example of an interactive music system made from crafting with digital technology [Torre et al., 2016]. Crafting may go in line with creative practices of music that require specific implementations as a form of human expression.

However, crafting has also its own limits. Even if a source of inspiration for many humans, crafting's sole relying on practice over theory may sometimes induce a lack of transmission [Jorda, 2005]. If not documented with care, many crafted interactive music systems—and the music practices and repertoire that go with them—may not subsist to our social history [Herbert, 2012]. The issue of interactive music system archiving remains contemporary, as today's fast evolution of computer hardware and software programming languages adds complexity to documentation [Risset, 2004].

There exists a continuum of approaches to the design of interactive music systems, ranging from bottom-up engineering sciences to top-down crafting practice. Careful balancing between the two extremities may be required to avoid respective pitfalls and fairly benefit humans.

2.1.3 Interaction Techniques

Now that we provided an overview of the approaches incumbent to their design, we are able to focus on the techniques at stake in interactive music systems. We propose to review them from the human-centred perspective of musical interaction, that is, from the new actions they provide humans with in relation to music practice [Bevilacqua and Schnell, 2016]. Relying on Beaudouin-Lafon's classification of human-machine interaction paradigms [Beaudouin-Lafon, 2004], we propose to classify interaction techniques in three categories: (1) instrumental control, (2) automatic generation, and (3) collective interaction.

Instrumental Control

Instrumental control originates from acoustic instruments, which essentially provided humans with one-gesture-to-one-event relationships [Jorda, 2005]. For example, percussive instruments enable humans to trigger short sounds by hitting some kind of surface with their body or some kind of object. As a widespread technique for music practice, instrumental control has naturally extended from physical to digital music instruments. The MIDI protocol, created in the 1980s, proposed a communication

standard for digital instruments that imitates the note-based control of acoustic instruments [Loy, 1985]. This notably enabled the building of tools relying on notation to create music, which in turn led to the birth of a music practice called computer-assisted composition [Assayag et al., 1999].

Gestural controllers have been developed to improve instrumental control in digital music systems, to go beyond mouses and keyboards by extending humans' physical abilities to produce sound. While keyboards were largely used following the MIDI protocol standardization, sensors provided new perspectives for instrumental control by fully redefining gesture-sound relationships in music systems [Miranda and Wanderley, 2006]. The process of "mapping", which specifies the relationships between sensor and sound parameters by means of programmation, received a lot of attention in the last decades [Hunt and Wanderley, 2002]. Unlike acoustic instruments, whose physics already define a mapping, digital instrument requires the relying on different strategies to design mappings—*e.g.*, one-to-one, one-to-many, many-to-one, or many-to-many [Jorda, 2005]. The OSC protocol eased the creation of mappings by extending MIDI to fully-generic communication between music systems [Wright, 2005].

Automatic Generation

Automatic generation refers to the ability of interactive music systems to create sounds, or sequences of sounds, based on anthropomorphic means of communication. It can be a desired feature of interactive music systems for music practices that take algorithmic composition as a main concept, or for those that may benefit from a certain degree of delegation to the machine.

While automatic generation remains highly characteristic of digital music systems, it was already a desired feature for many musicians centuries ago. For example, Mozart's *Musikalisches Wurfelspiel* ("Dice Music"), a musical game which "involved assembling a number of small musical fragments, and combining them by chance, piecing together a new piece from randomly chosen part", was one of the first algorithmic composition systems to date [Alpern, 1995]. Many computer-based interactive music systems made use of similar rules to provide composers with automatic generation. Lejaren Hiller is often reported to be the first composer to ever use a computer-generated score for a musical piece [Hiller and Isaacson, 1957], later followed, among others, by Iannis Xenakis and his "stochastic" compositions [Xenakis, 1992]. While reviewing artistic practices on automatic generation remains out of the scope of this thesis, we wish to stress that the notion of "artificial intelligence" was linked with music way before the advent of machine learning. Specifically, the computer music community introduced the term "intelligent instrument" [Spiegel, 1987] to describe interactive music systems that automate the control of music in some range [Chadabe, 1977]. Machine learning will be the matter of Section 2.2.

Alternatively, automatic generation can also be a way to *delegate* superfluous tasks to the machine, in a way similar to a human partner. For example, Cartwright et al. developed an interactive music system that enable querying of sounds using vocalization as human input [Cartwright and Pardo, 2015], thus enabling humans to focus on creation over search. Also, the delegation of tasks to the machine may enable humans to be more productive in their musical practice. For example, Deruty proposed

tools to automate the process of mixing to speed up music production [Deruty, 2016]. Improving music production may be useful for music practices that are connected to the industry.

Collective Interaction

Collective interaction refers to the ability of interactive music systems to facilitate interaction between humans. It is a promising avenue for music practices that focus on situated interaction, as well as creative collaboration, between humans.

Collective interaction is a central feature of many music practices situated in location. Whether it be through composition, performance, or listening, music is tied with social bonds that does not necessarily require an interactive music system to be mediated [Leman, 2008]. Recently, Matuszewski et al. investigated the use of mobile phones to create new interactions between humans situated in location [Matuszewski et al., 2019]. Networking technologies enable new forms of communication between professional musicians and the public, opening new prospects for collective interaction in a context of music pedagogy [Bell and Matuszewski, 2017].

Collective interaction may also be useful to manage collaboration in music creation. For example, shared multimodal interfaces enabled the creation of new collective music practices through the creation of [Jorda, 2005]. Cartwright et al. created a web interface that allows the crowdsourcing of semantic descriptors in relation to sound [Cartwright and Pardo, 2013]. Live coding is another example of collaborative music practice, where humans may collectively write lines of code to generate sounds and sequences of sound, possibly relying on Web frameworks [McKinney, 2014]. From a pedagogical point of view, web platforms enabled collaboration between musicians from all across the world by supporting the streaming and learning of music practices [Ruthmann, 2007].

Interactive music systems may be analysed as tools, partners, or mediums, for instrumental control, automatic generation, and collective interaction with music. Designing interactive music systems consists in blending these three techniques in careful accordance with situated music practices to support new forms of human expression.

2.2 Machine Learning

Machine learning has become one of the most successful techniques of Computer Science in the 21st century. Unsurprisingly, interactive music systems have witnessed a rise in applications of machine learning techniques. In this section, we review how machine learning was applied to the design of interactive music systems. We start by giving a formal definition of the notions of technique, data, model, and algorithm, that are central to machine learning. We then give an overview of the applications of machine learning to music, using a similar classification than in Section 2.1.3. Finally, we describe the research methods that have been employed to approach machine learning applied to music.

2.2.1 Formal Definitions

Machine learning studies the design of computational algorithms capable of performing specific tasks by autonomously learning from data. In this section, we provide formal definitions of the notions of (1) technique, (2) data, (3) model, and (4) algorithm, which are central to understanding applications of machine learning to music.

Technique

Machine learning defines different *techniques* that enables to perform tasks in an autonomous manner. Typical tasks are classification and regression—*e.g.,* respectively predicting discrete or continuous output for a new given input. Other typical tasks relate to optimal decision-making with a long-term reward trade-off.

The definition of a technique is made through the formalization of a learning problem in relation to a goal. In *supervised learning*, the goal is to learn a function that maps inputs to outputs based on example pairs of input-output [Bishop, 2006]. In *unsupervised learning*, the goal is to learn a similar function based on example input only [Murphy, 2012]. In *active learning*, the goal is similar to that in supervised and unsupervised learning, instead that learning occurs in interaction with a human labeller [Settles, 2010]. In *reinforcement learning*, the goal is to learn to take actions in an environment so as to maximize some notion of reward [Sutton and Barto, 2011]. Other learning problems exist, but remain out of the scope of the present thesis.

Data

Machine learning relies on *data* to perform the tasks defined above. Data is a set of values supposed to represent some kind of object with respect to qualitative or quantitative variables. For example, visual images, which physically emerge from complex sums of light waves, are often represented by a finite set of pixel data, the latter being made of red, green, and blue digital values. Data can have many roles in relation to machine learning techniques. *Example data* allows to perform supervised, unsupervised, and active learning—*e.g.,* respectively, example input-output, input only, and output only [Bishop, 2006, Murphy, 2012]. *Reward data* allows to perform reinforcement learning—*e.g.,* positive or negative reward for a given action [Sutton and Barto, 2011].

Specific *features* of data may be required to support learning. Feature selection is often performed to process data in a relevant way with respect to the task being considered. Based on this, larger *data sets* may be built to perform supervised, unsupervised, and active learning. Researchers have created standardized data sets to provide the machine learning community a common baseline to advance the building of learning algorithms [Bertin-Mahieux et al., 2011]. In the case of reinforcement learning, large amounts of data are synthetically created by means of a reward function. Many techniques for data processing are studied in the machine learning field, but remain out of the scope of this chapter.

Model

Models are mathematical functions capable of handling data to perform the machine learning task. Some models rely on relatively simple architectures, providing concrete representations of data—*e.g.,* the Gaussian distributions, for example used in supervised and unsupervised learning. Some models rely on complex architectures, providing multiple layers of abstractions—*e.g.,* the neural network models used in deep learning. Other models are sequential, that is, they provide a sense of the temporal process to be modelled—*e.g.,* the Markov decision process used in reinforcement learning.

In addition to the task being performed, models are sometimes used to provide qualitative insight on data. These data-driven models may be used as a complement, or an alternative, to mechanical models often used in science to describe the world. As such, the models' formalism may be carefully chosen depending on the type of data to be studied. Hybrid models may take advantage of both data-driven and mechanistic approaches [Alvarez et al., 2009], but remain out of the scope of this thesis.

Algorithm

Algorithms are a set of rules used for the optimization of the model's parameter in relation to a data set. Optimization is responsible for the *learning* aspect of machine learning, and is what makes machine learning different from statistics as a field [Bishop, 2006]. The algorithm's parameters are called hyperparameters, in the sense that they support the learning of the model's parameters. Some algorithms require large amount of data to be able to learn a model—*e.g.,* deep learning algorithms [Goodfellow et al., 2016]—but as a counterpart, may automatically extract features from data. Others are designed with the requirement of learning from small amounts of data—*e.g.,* shallow learning algorithms—, but in turn are highly sensitive to the choice of data features.

Importantly, learning can be performed *offline*, as is the case in most supervised and unsupervised learning—*e.g.,* training over a fixed data set, then performing the task without learning anymore. Or, it can be performed *online*, as is the case in most active and reinforcement learning—*e.g.,* training in an incremental way as new data feeds into the algorithm. The choice between offline and online learning algorithms depends on the task to be performed. Computational descriptions of these algorithms remain out of the scope of this chapter.

Machine learning defines a set of algorithms capable of autonomously learning a model from example data sets. The learned model may either be used to analyse the given data set, or to carry a task on newly-seen data.

2.2.2 Musical Applications

Now that we defined the notions that are central to machine learning, we are able to focus on its musical applications. Similar to Section 2.1.3, we propose to review them by relying on Beaudouin-Lafon's classification of human-machine interaction

paradigms [Beaudouin-Lafon, 2004]: we propose the following classification for interactive music systems: (1) audio analysis, (2) music generation, and (3) human-machine improvization.

Music Information Retrieval: Machine Learning as Tool

Music information retrieval uses machine learning to automate certain tasks related to music data [Downie, 2003]. For example, track separation aims at automatically extracting audio tracks of each musical instrument from a musical recording. This may be formulated as a supervised learning problem, where example data are made of input audio signals—*e.g,* a waveform, or a spectrogram—, along with their corresponding instrument label as output—*e.g,* bass. Another application is recommender systems, which aims at automatically creating music playlists by predicting a listener's musical tastes [Ricci et al., 2011]. This may be formulated as a reinforcement learning problem, where the system would learn to provide a listener with optimal rankings of songs in an online manner, based on reward given by other listeners and the listener herself.

The goal of music information retrieval is to create tools that may facilitate human interaction with music. The automation of certain tasks may be useful for professional creatives that may require a focus on productivity in their work [Deruty, 2016]. The created tools may be useful for industrial applications that have a focus on music consumption [Wikström, 2013].

Artificial Creativity: Machine Learning as Partner

Artificial creativity uses the models built in machine learning to automatically generate new sequences of sounds [Boulanger-Lewandowski et al., 2012]. For example, style modelling employs machine learning to generate music that share stylistic similarities with other music [Dubnov et al., 2003]. This may be formulated as a supervised learning problem, where style may be computed from specific data features—*e.g.,* tempo, harmony—, and generation would be made from a distribution model recreating data that share similarities with the training set. Instead of automatically creating sequences of sounds, other approaches to artificial creativity consists in creating new sound spaces [Esling et al., 2018]. This may be formulated as an unsupervised learning problem, where sound spaces are first learned from unlabelled audio data sets, then leveraged from the architecture of the learned models.

The goal of artificial creativity is to create artificial partners that may inspire humans by generating new musical ideas. However, the human benefits of these artificial partners remain to be studied on a longer-term before claiming that it augments human creativity [Roberts et al., 2019]. For example, professional musicians may enjoy creating music, and as such, may prefer doing it all by themselves rather than relying on automatically-generated structures. Also, even if these partners may help music composition be more approachable by non-musicians [Huang et al., 2019], they may not enable embodied interaction with music.

Human-Machine Improvization: Machine Learning as Medium

Human-Machine Improvization uses machine learning to adapt music generation to human musical data in real-time. The goal is to create dynamic processes of interaction between humans and machines. For example, Assayag et al. created a series of musical improvizers able to play and adapt to a musician's style of playing by learning from it in real-time [Assayag et al., 2006, Nika et al., 2017]. Similarly, Pachet et al. created a musical looper which automatically adapts the generation of a musical accompaniment depending on what the musician is playing [Pachet et al., 2013]. In test studies with professional musicians, these machines were not perceived as emotionally intimidating, but rather as musically expressive [Lévy, 2013]. This positive embodiment of machines by humans may provide new opportunities to play (with) music.

The task of human-machine improvization may promote machine learning to the function of *reflexive medium*. Similarly to the video medium, which is "in constant movement, circulating between camera and monitor" [Spielmann, 2010], the machine learning medium may be in constant movement, circulating between human and music. Rather than static descriptions of reality, the computed models may be considered as diachronic forms, that is, forms that depict the transformative nature of reality [Bourganel et al., 2016]. online learning from real-time human data may enable this accounting for a dynamic nature of reality. All in all, the newly-created medium may encourage improvizational interaction between humans and machines, which, just as human-human improvization, may benefit human expression [Wright and Kanellopoulos, 2010].

Applications of machine learning to music have so far focused on music information retrieval, music generation, and human-machine improvization, to create new musical tools, partners, and mediums. The latter task opens new prospects to build interactive music systems that supports embodied forms of human expression.

2.2.3 On Research Methods

In this section, we give an overview of the research methods used to formalize machine learning and lead applications in music, in light with the design approaches of interactive music systems described in Section 2.1.2. We describe (1) the engineering sciences approach and (2) the thing-ification of music and human creativity at stake in machine learning.

The Engineering of Machine Learning

As a relatively recent field emerging from Computer Science and Mathematics, most applications of machine learning to music have been done from an engineering sciences perspective, that is, the solving of a technical problem. The evaluation of machine learning then consists in measuring the performance of the engineered model in relation to the technical problem. However, we argue that technical problems may not always be related to human problems in music practice. Furthermore, these hu-

man problems may not always be resolved with a technical solution. While engineering sciences may be necessary for the music information retrieval task—which is by essence quantitative—, it may not always suit the tasks of artificial creativity and human-machine improvization—which crucially pose qualitative issues regarding music practices and human interaction [Lévy, 2013]. Including humans in the design of machine learning may help balance machine engineering with human goals.

Engineering a model is a long and tedious research process. It involves the creation of a—possibly large—data set, and the optimization of parameters of a model in relation to this data set. The computational resources needed to fulfill these activities—the storage of data and the learning of optimal parameters—can become quite high compared to standard computers' performance. Also, the learning of optimal parameters on a given data set may not transfer to another data set. As a consequence and despite the real-world success of some applications, many machine learning applications remain not applicable in human situations because of this focusing on machine task engineering [Wagstaff, 2012].

The Thing-Ification of Music and Human Creativity

Many formalizations of music as a machine learning problem tends to consider music as a thing, not a practice. While data representations, learning algorithms, and models remain completely generic from a computational point of view, researchers remain mostly focused on Western music theory as a benchmark to engineer models—*e.g.,* melody, polyphony, accompaniment, counterpoint [Briot et al., 2017], or electronic dance music [Eigenfeldt and Pasquier, 2013]. Also, while effort is made toward including multicultural perspectives on machine learning [Serra, 2017], the creation of large data sets of music still tends to fix music as a static thing. Yet, as described in Section 2.1, music is a dynamic practice that is subject to cultural evolutions [Cross, 2001a].

The "thing-ification" of music may also direct the research of machine learning toward developing artificial creativity over human creativity. The loosely-defined concept of "artificial intelligence" nurtures the phantasm of creating machines that are capable of human-level creativity [Colton et al., 2012]. While we do not claim that all machine learning researchers share these views on "artificial intelligence", we argue that they do not always clearly take a stance on this issue. Of course, individual researchers may not become experts in all other disciplines involved in the notion of "artificial intelligence" within their available resources [Mackay and Fayard, 1997]. Yet, as many of their predecessors [Wiener, 1988], we believe that today's computer scientists should actively include their work as part of our society [Latour, 1987]. Interdisciplinary collaborations may enable new reformulations of complex scientific subjects [Mackay and Fayard, 1997], including machine learning.

Machine learning remains mostly studied from an engineering sciences perspective, and applied to music without a clear accounting of its practical aspects. Key machine learning notions must be reformulated through interdisciplinary collaboration to lead situated applications in music and benefit human expression.

2.3 Human-Centred Machine Learning

Human-centred machine learning is an interdisciplinary area of work that aims at rethinking machine learning research and systems in terms of human goals. In this section, we motivate human-centred machine learning for the design of interactive music systems. We give an overview of its research approaches, which investigate machine learning techniques through the lens of human-centred methods. We finally review musical applications of human-centred machine learning.

2.3.1 Motivation

Human-centred machine learning includes researchers and practitioners from the fields of (1) human-computer interaction and (2) interactive machine learning. In this section, we motivate a human-centred approach to the design of machine learning in interactive music systems by introducing these two areas of work.

Human-Computer Interaction

The field of Human-Computer Interaction (HCI) researches the design and use of new computer technology by humans. Historically originating from the Engineering Sciences and Computer Science with the aim of designing efficient human-computer interfaces, the field has progressively evolved to study the interaction between humans and computers [Beaudouin-Lafon, 2004]. Specifically, the last decade saw the rise of *embodied interaction* as part of the third wave of HCI, which included aspects such as experience, emotion, or bodily interaction in the design of computer technology [Dourish, 2004].

Over the last two decades, the HCI community expanded to study the design of creativity support tools [Shneiderman, 1999]. The focus of this topic is to understand the creative process of human users to design interactions with computers that facilitate creativity [Resnick et al., 2005]. Music was in this sense examined as a use case to design technology that is adapted to composers [Garcia, 2014], or reciprocally, as an inspiring modality to design new interactions with technology [Ghomi, 2012]. Recently, creativity researchers called for a move to interdisciplinary collaborations [Frich et al., 2018].

Interactive Machine Learning

The field of interactive machine learning studies the use of machine learning systems by humans. It originated a decade ago, from the assessment by machine learning researchers that a lot of expert knowledge was required to actually make a machine learning system work [Domingos, 2012]. Researchers wanted to give "power to the people" by building machine learning systems able to account for their human users [Amershi et al., 2014]. For example, studies shown that humans often interacted with artificial agents in a way that may not suit the underneath machine learning implementation [Thomaz and Breazeal, 2008]. *Machine teaching* was introduced as conceptual framework that focus on empowering human users [Simard et al., 2017].

Computer music pioneered interactive approaches to machine learning for creative applications, with real-time programming environments such as Max/MSP enabling exploration of models for sound synthesis [Lee and Wessel, 1992, Fels and Hinton, 1993, Bevilacqua et al., 2005]. More recently, interactive machine learning was applied to the design of musical design tools [Fiebrink, 2011], with two complementary human- and machine-centred advantages. First, it enabled to raise relevant aspects of human creativity in relation to music practice, such as surprise, unexpectedness, and exploration [Fiebrink et al., 2010]. Second, it supported the building of better machine learning software used by real-world human users, enabling the teaching of machine learning to musicians and non-musicians [Fiebrink, 2019].

Human-centred machine learning blends approaches from the field of Human-Computer Interaction and Interactive Machine Learning. Applications may potentially benefit music practice by taking into account various dimensions of human creativity in the design of machine learning systems.

2.3.2 Research Approach

This section describes the research approaches in human-centred machine learning, from (1) interactive, open-ended learning to (2) human-centred, scientific methods.

Interactive, Open-Ended Learning

Human-centred machine learning is essentially concerned with the *interactive learning* of a model from *user-provided data* [Fails and Olsen Jr, 2003]. We use the word user here to underline the fact that data may not be harvested by engineers or automatically by implicit measures, but actively created by the human users of the machine learning system. As such, users may iteratively add, delete, or modify example input-output data in supervised learning to tacitly engineer a new model over a customized data set. They may also experiment with various algorithm hyperparameters over a unique data set, for example adjusting error preferences in an unsupervised learning-based classifier [Kapoor et al., 2010].

Once they trained the model, they may evaluate it in relation to the task it has been applied to. Rather than the quantitative, goal-oriented criteria used by engineers, human users may use *qualitative, open-ended criteria* to evaluate the working of a model. For example, Fiebrink et al. showed that musicians may evaluate a gesture classifier from the sonic output that it may produce, rather than on the actual recognition rate produced by the model [Fiebrink et al., 2011]. *Experiential aspects* of interactive learning may also be important for user evaluation, as several applications in the field of robotics showed—*e.g.*, not wanting to spend too much time giving reward in reinforcement learning [Thomaz and Breazeal, 2008], or not wanting to be asked too many queries in active learning [Cakmak et al., 2010].

Human-Centred, Scientific Methods

Human-centred machine learning essentially relies on methods from the field of Human-Computer Interaction to deal with user evaluation. *Design-oriented* methods focus on individual construction and implementation to create innovative interactive systems [Consolvo et al., 2009]. *Experimental methods* assess a clearly-defined hypothesis linking a given machine learning parameter with a user experience aspect using a controlled setup [Lazar et al., 2017]. *User-centred design* methods triangulate between theory, design, and observation to understand the needs of human users, design new technology, and put it into use by the same users [Mackay and Fayard, 1997]. *Participatory design* methods put an emphasis on all stakeholders of a project in the design process of a new technology [Schuler and Namioka, 1993]

While HCI includes aspects of crafting, it cannot be considered as a pure design discipline in itself, "because it strives to independently verify design decisions and processes, and borrows many values from scientists" [Mackay and Fayard, 1997]. *Prototypes*—a concrete representation of part or all of an interactive system—are central to these design processes and may be considered as *design artifacts* [Beaudouin-Lafon and Mackay, 2009], as well as to "the design goal of inspiring users and researchers to think about technology" [Hutchinson et al., 2003]. Yet, HCI aims at creating and revising interaction frameworks between users and technology [Mackay and Fayard, 1997]. Thus, most HCI methods share the scientific goal of evaluating the benefits of a new technology for their users to contribute to a model of interaction [Hutchinson et al., 2003].

Human-centred machine learning studies the interactive and possibly open-ended uses of machine learning models by humans. It applies scientific methods to iteratively design prototypes, evaluate human interaction, and build frameworks for human-machine interaction.

2.3.3 Example Musical Applications

In this section, we describe three example musical applications of human-centred machine learning. We present (1) The Wekinator, (2) Grab-and-play, and (3) CoMo, applying machine learning to the practice of motion-sound mapping.

The Wekinator

The Wekinator is a software developed by Rebecca Fiebrink and colleagues at Princeton University, which uses interactive machine learning to let humans map any motion sensor to any sound synthesis engine [Fiebrink et al., 2009]. Specifically, it uses supervised learning to create a model that links motion input to sound synthesis output. This model is learned from examples of gestures and sounds that human users may record. They may interactively add, delete, or modify gesture-sound examples, train a model, and experiment with the newly-created model to build their customized model. Typical models include gesture classifiers—built with k-nearest neighbour algorithm—, and sound regression—built with neural network algorithm.

The software was developed through a user-centered design process with many expert musicians implied. It was then applied in many real-world contexts of music practice, such as music composition [Fiebrink et al., 2010], instrument design and performance [Fiebrink et al., 2011], but also for music pedagogy [Morris and Fiebrink, 2013, Fiebrink, 2019]. In all cases, humans valued the interactive supervised learning workflow created by the Wekinator. Specifically, they appreciated the surprise and unexpectedness enabled by the system, which fostered their creative process. As such, interactive supervised learning may be seen as a similar task to human-machine improvization, where the machine serves as an expressive partner for creation [Fiebrink, 2017].

Many other systems for motion-sound mapping were developed using a similar interactive supervised learning workflow, but different algorithms enabling other tasks—*e.g.,* gesture following [Bevilacqua et al., 2009], gesture variation estimation [Caramiaux et al., 2015], or probabilistic modelling [Françoise et al., 2014].

Grab-and-play

Grab-and-play is an extension to the Wekinator that I developed during a pre-doctoral research year at Goldsmiths University of London, under the supervision of Rebecca Fiebrink [Scurto, 2016]. It uses interactive machine learning to enable human users to literally grab a controller and turn it into a new, playable musical instrument almost instantaneously. Specifically, it uses supervised learning to create a motion-sound model from a very small amount of data communicating some minimal, soft design constraints—namely, the way the user might want to move while playing. As a data-driven extension, Grab-and-play leverages the same learning algorithms than the Wekinator.

The system was developed through a two-step user-centred design process with composers on the one hand, and disadvantaged children on the other hand. It was then respectively applied to instrument design [Scurto et al., 2016] and music therapy [Thompson et al., 2019]. The rapid prototyping enabled by Grab-and-play favoured ideation in the creative process of composers and performers, as well as a better focus on the expressive needs of children for music therapists. The diachronic form supported by Grab-and-play makes it similar to a medium, as depicted by the task of human-machine improvization.

CoMo

CoMo is a prototype web application developed by the ISMM group of IRCAM, using interactive machine learning to support movement interaction with sound, specifically targeting collective interaction using with phones [Matuszewski et al., 2018]. The functioning is similar to the interactive supervised learning workflow of Wekinator—human users recording pairs of gesture-sound examples to build custom models—, except that it works at a distributed scale. Specifically, it enables one human user to build a custom gesture-sound model, and to share it with other human users using the web architecture. Typical models include gesture classifiers—built with gaussian mixture models—and followers—built with hidden markov models.

The application benefits from extensive participatory design processes implying many performers and movement practitioners and is still under development. It was applied to real-world situations in dance and pedagogy [Voillot et al., 2019]. The space for exploration created by the application benefited the creative practice of dancers, as well as improved playfulness in pedagogical activities. As a machine learning-based distributed application, CoMo also shares the qualities of a medium.

Human-centred machine learning applied to music creates reflexive mediums enabling new forms of human expression situated in music practices. For now, only supervised learning and motion-sound mapping have been investigated as machine task and human practice.

2.4 Summary

This chapter reviewed the context and related work on the design of machine learning for interactive music systems. We described the cultural and individual aspects of music practices, which have to be taken into account when designing systems for human expression and embodied interaction. We discussed how crafting and engineering may be two complementary approaches to designing the related technical components. We described how the final interactive music systems may be analysed as providing humans with instrumental control, automatic generation, and collective interaction with music to foster creativity.

We then defined machine learning as a set of algorithmic techniques that enable to automatically learn a model between input and output from example data only. We reviewed musical applications of machine learning and their enabling of new forms of musical expression, encompassing music information retrieval, music generation, and human-machine improvisation. We discussed how interdisciplinary collaboration may be necessary to situate the research of machine learning in existing music practices and support the design of novel interactive music systems.

Finally, we introduced human-centred machine learning as an emerging field studying Interactive Machine Learning from a Human-Computer Interaction perspective. We analysed how the field jointly considers qualitative and quantitative aspects of human interaction in the design and evaluation of machine learning. We described three example musical applications of human-centred machine learning to motivate our choice to adopt this research and design approach for this thesis.

3

Designing With Unsupervised Learning For Motion-Sound Mapping

This chapter presents the first study led in the thesis. We decided to focus on the musical task of motion sound mapping to build on previous applications of human-centred machine learning to music (see Section 2.3.3 for a review). Our wish was to generate ideas inside our research group before eventually involving human participants in the design process. As such, our method for this study can be described as design-oriented, where knowledge on both machine learning and the musical task would only be attained attained through the implementation of a model prototype [Fallman, 2003]. Figure 3.1 schematizes the three steps resulting from our use of the method in the study; we describe them below.

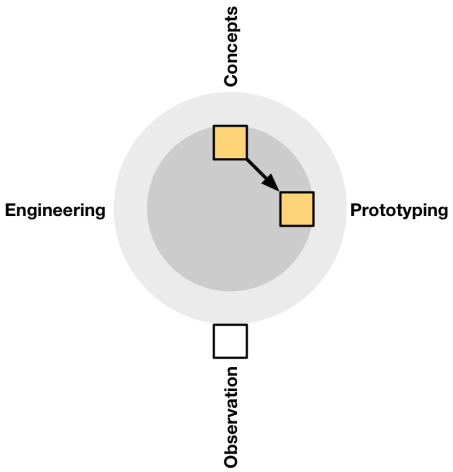


Figure 3.1: The design-oriented method for our first study.

Section 3.1 describes **concepts** on motion-sound mapping, covering both human practices and machine systems. Section 3.2 describes the **prototyping** of our model prototype. Finally, Section 3.3 describes the **application** of the model to the design of a musical instrument, called *somasticks*, which we exhibited to expert motion practitioners in a workshop.

This work has been published and presented as paper and poster at the *17th International Conference on New Interfaces for Musical Expression (NIME 2017)* in Copenhagen, Denmark [Scurto et al., 2017]. The musical instrument was showcased at the *movA workshop* (March 2019) in Nantes, France.

3.1 Concepts

We start by reviewing the concepts that are relevant to this first study, describing the related work in the musical task of motion-sound mapping, as well as the machine learning technique that we proposed.

3.1.1 Musical Task: Motion-Sound Mapping

Situating the Musical Task

Motion-sound mapping is a music practice consisting of programming the relationships between an input device—typically, a motion sensor—and an output device—typically, a sound synthesis engine [Hunt and Wanderley, 2002]. It is a creative practice that has long been the specificity of crafters, performers, artists, scientists, and engineers from computer music communities—especially the NIME community.

Recently, motion-sound mapping has become increasingly accessible through the use of interactive machine learning. With these technologies, humans can build custom motion-sound mappings by physically demonstrating examples of gestures for given sounds—thus relying on corporeal knowledge instead of programming skills.

Human-Centred Approaches

Most interactive approaches to machine learning for designing motion-sound mappings have relied on a two-step, iterative design process (see figure 3.2) [Françoise et al., 2014]. In the first step, called training or *design* step, users perform gestures *along with* pre-defined sounds. In the second step, called *performance* step, users experiment with the newly-created mapping. For example, they can perform similar gestures to the ones they recorded during the design step in order to replay, or reenact, previously-selected sounds; or, they can perform new gestures in order to explore, and discover, new sonic forms. Users must then alternate several times between these two steps in order to succeed in building a subjectively-rewarding mapping.

Several user studies have proven that this iterative design process can support corporeal engagement with sound [Bevilacqua et al., 2009, Fiebrink et al., 2010]. However, recent works have raised a number of points yet to be improved [Scurto et al., 2016]. For example, some users may have difficulties in designing gestures and evermore to

fine-tune mapping. Importantly, Scurto et al. found that users might appreciate machine learning-based mappings that surprise and challenge them through continuous physical interaction [Scurto et al., 2016].

Research by Fiebrink et al. have shown that machine learning can support human-centred creative discoveries in musical motion-sound mapping design [Fiebrink et al., 2011]. For example, criteria such as unexpectedness and accessibility have been praised by computer musicians when composing an instrument [Fiebrink et al., 2010]. In this spirit, Scurto and Fiebrink proposed new methods for rapid mapping prototyping which shift users' focus from designing motion-sound relationships to the embodied exploration of relationships that have been generated partly by the computer [Scurto et al., 2016].

Machine-Centred Approaches

While most initial research focused on explicit motion-sound relationship programming [Hunt and Wanderley, 2002], using machine learning algorithms for mapping design have proven very promising in a musical context where notions of expressivity and generativity are of prime interest [Caramiaux and Tanaka, 2013].

In this context, several supervised algorithms have been studied, depending on the musical task users would like to achieve. For example, Bevilacqua et al. [Bevilacqua et al., 2009] focused on gesture following tasks and implemented a Hidden Markov Model to perform continuous tracking on users' gestural data. Fiebrink et al. investigated static mapping building using neural networks for regression tasks and several standard algorithms for classification tasks, such as k-nearest neighbors [Fiebrink et al., 2011]. Françoise et al. proposed four static and/or dynamic models able to perform both classification and regression tasks [Françoise et al., 2014]. Finally, Caramiaux et al. developed a system that recognizes gestures and adapts to performance variations [Caramiaux et al., 2015].

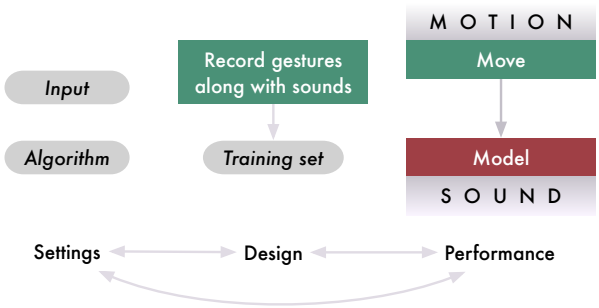


Figure 3.2: Interactive workflow for *Mapping-by-Demonstration*.

However, to our knowledge, most of these approaches remained focused on a two-step design process (see figure 3.2), where users alternate between demonstrating gestures along pre-recorded sounds (movement acted from the experience of listening to a sound) and interacting with newly-created mappings (movement acted as having an effect on sound). This iterative process might interrupt musical intentionality encod-

ing, which, as theorized by Leman, necessitates an active, action-oriented, corporeal engagement of humans with sound [Leman, 2008]. Interestingly, other computational approaches aimed at providing users with such continuous interactive flows, for example using dynamic mapping strategies [Momeni and Henry, 2006] or physics-based mappings [Schacher et al., 2014].

3.1.2 Proposed Technique: Unsupervised Learning

Inspired by such approaches and other interactive music systems [Jorda, 2005], we propose to reconsider mapping creation to bridge the gap between design and performance steps. We decide to focus on the machine learning category of unsupervised learning to perform the task of online clustering. Our wish was to allow mappings to adapt to users in real-time while generating sound, thus merging design and performance steps into one fully interactive experience.

Such a task differs from previous interactive supervised learning approaches: instead of demonstrating gestural examples that have been designed and labeled in a separate step, users physically interact with an adaptive model that constantly generates sound, depending on both previous and current user movement. Importantly, our system thus switches from current mapping-by-demonstration supervised paradigms (where user-provided pairs of gestures and sounds constitute a training set) to an unsupervised learning paradigm (where the training set consists in unlabeled gestural data). However, as we will see, users still have the possibility to consciously influence the learning by performing and correcting the system.

3.2 Prototyping

We describe the model prototype that we designed, in terms of interactive workflow and algorithmic implementation.

3.2.1 Algorithm: *Online Gaussian Mixture Model*

We implemented an online, unsupervised version of *Gaussian Mixture Model* (GMM). GMMs are very general and versatile probabilistic models for designing motion-sound relationships, providing with variables for both classification and regression at a relatively low computational cost [Françoise et al., 2014].

A GMM is a learning model that can perform soft clustering, which is identifying groups of similarity in gestural data and computing for a new data point $\mathbf{x}$ each probability that it belongs to each of these clusters. Here, clusters are modelled as Gaussian distributions $\mathcal{N}$, and the probability p of belonging to the overall model θ is given by:

$$p(\mathbf{x}|\theta) = \sum_{k=1}^{K} \pi_k \mathcal{N}(\mathbf{x}|\mu_k, \mathbf{\Sigma}_k) \tag{3.1}$$

There are four categories of parameters in GMM (see figure 3.3 and equation 3.1). The first one is the number of clusters K, which is the number of multivariate Gaus-

sian distributions used in the mixture model. These clusters can be used for classification purposes. Then, each Gaussian distribution has its own mean vector μ_k and covariance matrix $\mathbf{\Sigma}_k$, as well as its own weight π_k in the mixture. These parameters can be used for regression purposes. In a standard interactive supervised learning setup, such parameters are set and learnt offline from custom gesture-sound examples demonstrated by users. In our paradigm, the learning is online: Gaussian parameters would evolve in real-time as users supply the model with only gestural data, which support continuous action-perception workflow as specified in the previous section.

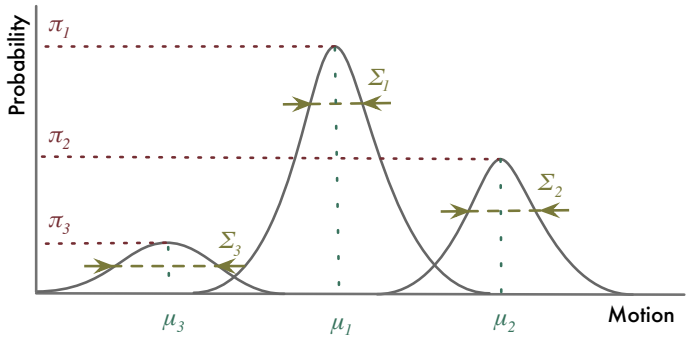


Figure 3.3: Graphical outline of a *Gaussian Mixture Model* with $K = 3$ cluster components for 1-dimensional motion data.

In such an online, unsupervised paradigm, we propose to add entropy $H = -\sum p(\mathbf{x}) \ln p(\mathbf{x})$ as a supplementary parameter for controlling sound synthesis. Entropy can be linked to the amount of information produced by some stochastic source of data [Liu, 2018]. In our case, a human that would always move in the same way would have a low entropy, as he or she would produce few "novel" information. On the other hand, a human that would have lots of variation in his or her moves would produce a higher entropy. Our idea is to report such a high-level measure of the model to use it as another modality for sound generation.

3.2.2 Human Data: Observation

Our wish is to enable users to design mappings in an online fashion, where design would be made possible through performance. We propose the following workflow, which is depicted in figure 3.4.

General workflow

Our system allows users to design machine learning-based motion-sound mappings while performing with them. More precisely, it enables online multidimensional adaptation to users input gestural space by continuously recording input data as the training set of a machine learning algorithm. Both design and performance steps are thus supported under the same motion flow. The modelling of the "internal structure" of users' gestural space can then drive sound synthesis in several manners, all of them

being characterized by direct, corporeal interaction with sound and personalized exploration of motion in relation to sound. We designed our system with a particular focus on reducing GUI actions taken in-between performances. One level of interaction with machine learning still remains available to users: similarly to previous mapping-by-demonstration tools, the "setting" step allows for configuring a minimal set of learning parameters as well as input parameters (as described in section 3.2.1).

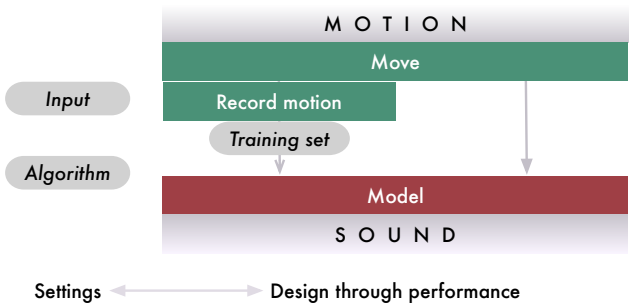


Figure 3.4: Interactive workflow for *Design Through Performance*.

Interaction modes

From this definition, we designed two interaction modes based on different memory processes (see Video for an overview). The main concept is to allow users to design parts of their input space through the metaphor of temporal persistence, where "occupation time" (as an "accumulation process") is central to the creation of the training set. There are several other ways to interact online with a machine learning algorithm through its training set: we will discuss it in section 3.4.

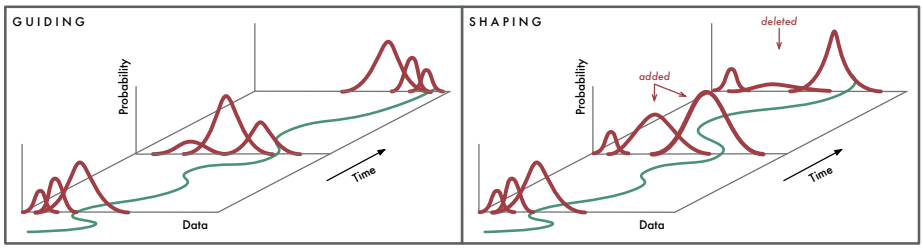


Figure 3.5: Graphical outline of our two interaction modes. Probability clusters are sampled at 3 discrete times for 1-dimensional gestural data. On left, clusters continuously evolve as users' gestural data is recorded to the training set with a sliding window. On right, users continuously modify clusters' parameters as they successively add and delete gestural data to the training set.

Guiding The guiding mode (figure 3.5, left) consists in having users adding gestural data with a sliding temporal window to the training set during the interaction. It

can be seen as an interactive music system with a constant-size memory, where users could directly and physically explore sound spaces in order to foster creative discoveries. It allows mappings to evolve continuously, focusing in or out of some spaces in users' gestural input space in real-time following abstract embodied specifications of users. A typical situation would involve the creation of clusters in a relatively small area of the input space by having users stay in this part of the input space, then its real-time evolution (or guiding) by moving in larger areas of the input space. This personalized interaction relies on an finite memory process where old data would be continuously replaced from the training set by new data.

Shaping The shaping mode (figure 3.5, right) consists in having users interactively adding and/or deleting gestural data to the training set during the interaction. It can be seen as a continuous extension of previous interactive machine learning systems, where users could delete and re-add a previously-recorded example in a design step by clicking on a button in a design step, then see the effect in a performance step. Here, users can add new examples and delete old ones by (re-)demonstrating them, while hearing the sonic consequences in real-time. Like using a pencil with eraser, this would allow rapid, custom, and fine-tuned modification of mappings. A typical situation would involve the creation of a new cluster for a new gesture, then its modification (or shaping) by adding or deleting variations of this gesture in the recorded data. This personalized interaction relies on an (almost-)infinite memory process where the training set would grow as users successively supply the system with data.

3.2.3 Implementation: Patch

We implemented a prototypical version of our system as a Max/MSP patch[1] that makes an extensive use of *XMM* library for learning GMMs [Françoise et al., 2014] and *MuBu* objects for storing and editing data [Schnell et al., 2009]. The tool's GUI provides users with different capabilities (see figure 3.6):

- Connect any kind of gestural input device, provided its data is sent as an OSC message.
- Experiment with different kinds of sound synthesis module, provided they receive OSC messages.
- Modify the training set physically either by adding, deleting, or window streaming gestural data.
- Define the length of the recording window.
- Define the number of Gaussian components in the GMM.

Currently, our tool supports online learning by training and running a GMM at a sufficiently high rate (every 100 ms) so that it remains perceptually convincing in an action-perception workflow [Jorda, 2005]. Gestural data is either stored incrementally or replaced dynamically by making use of overdub and append messages of the MuBu

[1] http://github.com/hugoscurto/OnlineGMM

container. The “delete" action is made possible by identifying and deleting the first nearest neighbour of user live input in the training database. Other implementations and user interfaces could be imagined; we will discuss them in Section 3.4.

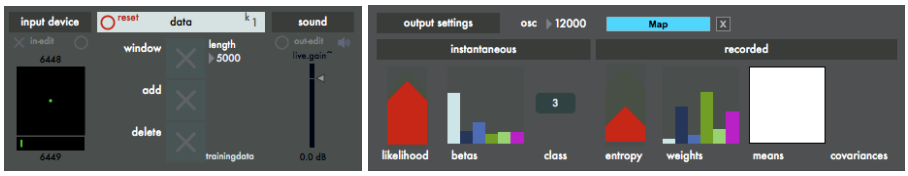


Figure 3.6: User interface of our model prototype. On left: Main window allowing recording gestural data following different interaction modes. On right: Output model parameters.

3.3 Application: *somasticks*

We applied our model prototype to the design of a musical instrument, named *“somasticks”*. We led a demonstration at the *movA* workshop, in Nantes, France, where expert performers and movement practitioners interacted with the instrument.

3.3.1 Description

somasticks are augmented drumsticks that seek to emphasize the somatic side of drumming practice. Contrary to standard drumsticks, *somasticks* do not need to hit any physical objects to produce sound, but rather leverages on embodied listening to drive musical performance. Specifically, they may be continuously waved in the air to trigger recorded drum sounds, and explore various playing modes in reaction to the internal sensations that the sounds may produce within the player’s body.

somasticks combine unsupervised learning with hardware elements. We used real drumsticks to create gestural affordances that are naturally related to drumming practice. We embedded the sticks with wireless sensors to feed the Online *Gaussian Mixture Model* with drumming motion data. We finally leveraged the online behaviour of the unsupervised learning model to design interactive drumming sound processes.

3.3.2 Instrument Design

Hardware

We designed hardware for *somasticks* in collaboration with IRCAM’s Emmanuel Fléty and Arnaud Recher from the PIP group, and Djellal Chalabi from the S3AM group.

Object We used wooden tip drumsticks as physical controllers for *somasticks*. At the time of prototyping, we used Vic Firth AJ5 American Jazz model, which are wooden tip, lightweight drumsticks designed for standard use. The sticks are 406 millimeters long, 12 millimeters wide, and weigh 36 grams. Any other drumsticks may be used in lieu of these in a customisation setting.

Sensor We equipped each of the sticks with a three-dimensional Inertial Measurement Unit developed in the ISMM team, called R-IoT[2]. In addition to the embedded accelerometers, gyroscopes and magnetometers used for motion sensing, we linked three button-like elements to each R-IoT. One piezoelectric sensor allows for the continuous control of one given parameter; two push-buttons allows for discrete triggering of given processes.

Sensor Support We designed a support for linking the sensors to the sticks (see Fig. 3.7). A clippable part allows to fix the R-IoT and its button elements to the body of the stick, using a screw to adapt to different stick diameters. A sliding part allows to set the position of the R-IoT battery to adjust the center of mass of the augmented stick. The model for the support is 3D-printable and thus allows for potential reproduction.

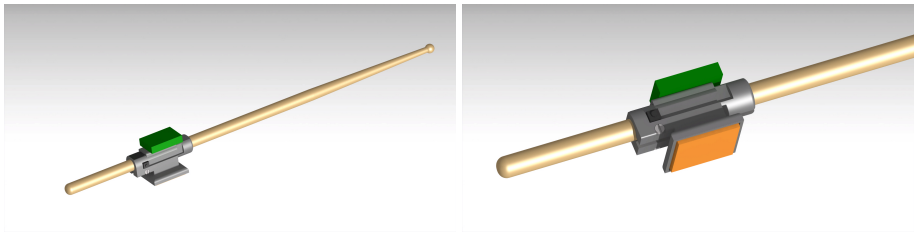


Figure 3.7: Rendering for the 3D-printed support of the *somasticks*. Left: clippable part. Right: sliding part.

Software

We designed a Max/MSP patch for *somasticks*, which converts the sticks' data into drum sounds.

Motion analysis We used online wavelet analysis, developed by Jules Françoise and the ISMM team in the frame of the SkAT-VG[3] European project, to extract motion features and inject them in the unsupervised learning algorithm. Wavelet analysis allows to measure the frequencies in a given movement, which is relevant as movements in drumming practice often carry periodicity. We performed one wavelet analysis per stick over its R-IoT accelerometer data, and took the maximum amplitude bin as the main frequency feature for a given stick.

Sound synthesis We used concatenative synthesis, developed by Diemo Schwarz and the ISMM team, to generate new sound patterns from recorded sounds [Schwarz et al., 2006]. Concatenative synthesis allows to classify sound samples depending on their spectral content, which is useful for interacting with large corpuses of sound. We created six sound corpuses related to six percussion elements commonly at stake in standard drum kits: bass drum (65), snare drum (77), rack toms (82), hi-hat (133),

[2] http://ismm.ircam.fr/riot/
[3] http://www.skatvg.eu/

crash (50), ride (47). The wide number of recorded samples provides a variety of playing modes on the drum kit elements—*e.g.*, from soft to hard hitting—, which is easily analysed and manageable using concatenative synthesis.

Interactive mapping We used the Online *Gaussian Mixture Model* to design the interactive behaviour of the *somasticks* (see Fig. 3.8). We used the guiding mode with a recording window of 10 seconds. The two sticks' maximum amplitude wavelet bins serve as two-dimensional data $\mathbf{x}$ for the model. Each sound corpus is mapped to one gaussian component, making our model contain $K = 6$ cluster components.

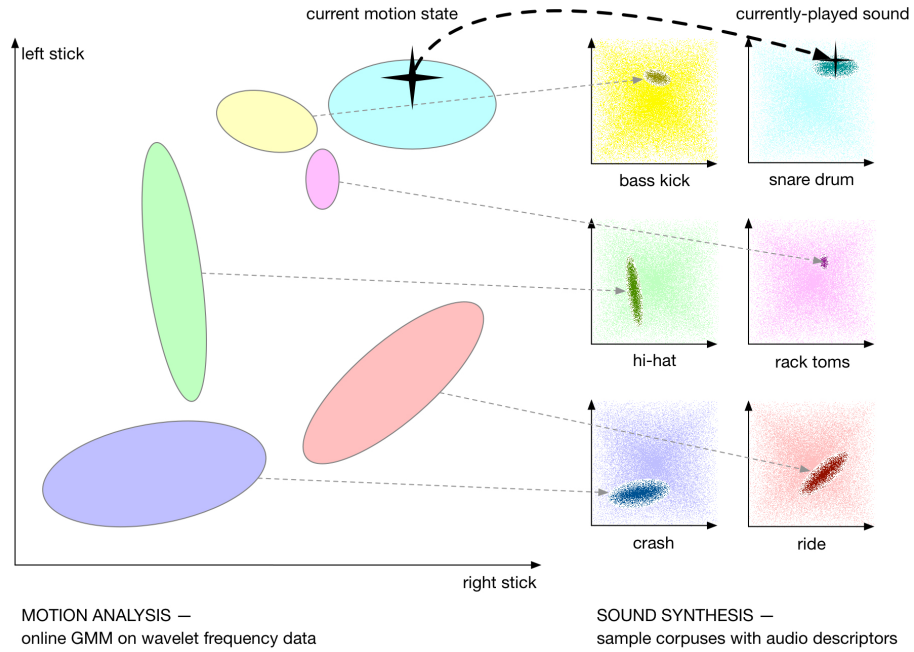


Figure 3.8: Graphical outline of the mapping used in *somasticks*. Each of the $K = 6$ clusters of the model (left, in colors) corresponds to one sound corpus (right, with corresponding colors). The means μ and covariances Σ of the clusters (ellipses) are mapped to the descriptor spaces of their respective corpuses, defining a region from which samples would be played. All clusters evolve online based on the performer's motion. In the current figure, the performer is in a motion state corresponding to the instantaneous class of snare drums: *somasticks* will thus play snare drum sounds, with synthesis qualities defined by both instantaneous and recorded model parameters.

Instantaneous model parameters handle low-level behaviour of the mapping. The model probability p sets the period at which new samples are played: the higher the probability is, the faster the samples would be played. The instantaneous class K defines the sound corpus from which samples would be played. The Gaussian probabilities $\mathcal{N}$ set the respective gains at which samples from the K classes would be played. Our intention was that the performer would have to produce stable quality movements in terms of frequency to be able to control the *somasticks*' sound output.

Recorded model parameters handle high-level behaviour of the mapping. The weights π set a random variation in the temporal moment where a sample would be played: the more weighted a cluster is the more regular the playing of its samples would be. The means μ and covariances Σ define the region of the sound corpus in which samples would be searched, using a k-nearest neighbor algorithm. We used two descriptors from concatenative synthesis (FrequencyMean and LoudnessMean) to set the searched sound space, and scaled them over the clusters' mean values. We computed the two eigen values from the covariance matrix and used their mean value as a mean radius for the k-NN. Entropy was tested in a series of preliminary experiments, but was not used in the latest version of the instrument due to issues in implementation.

Additionnal controls Lastly, the buttons enable direct, non-autonomous control of sound synthesis. The two piezoelectric sensors respectively control the period range over which samples are played using p, and a resampling variation for the samples played. The four push-buttons respectively control: activation of the gain scales set by $\mathcal{N}$, setting of a short, absolute duration (10 ms) for all played samples, activation of online training, and activation of sound.

3.3.3 Demonstration in Workshop

We demonstrated *somasticks* at the *movA workshop*, in Nantes, France, during which six expert movement practitioners tried and interacted with them. Our aim was to harvest reactions on the interactive behaviour of the sticks, understanding how its online implementation may be felt by different movement experts.



Figure 3.9: Picture of a performer playing with *somasticks* in the *movA workshop*.

Procedure

Participants were asked to spend time interacting with the *somasticks*. They were told to follow the following feedback loop: listen to the produced drum sounds, focus on their own corporeal sensations, and move freely with the sticks and their interactive sonic feedback. We only told them that the sticks were responsive to their movements' dynamics, and did not explain the sticks's internal functioning. Participants spent between 5 and 20 minutes exploring the system.

Reactions

Discovering the sticks All participants started interaction by exploring the interactive behaviour of the sticks. While this exploration was meant to understand the functioning of the software mapping, it was hugely influenced by the hardware affordances provided by the sticks: *"the sticks, really, define their very own geometry, so I focused a lot on the trajectories defined by the lines [of the sticks]"*, one participant said.

Reflecting on the sticks

The need for time After discovering the sticks, four of the six total participants took time to reflect on their functioning. One of them, who was a dancer, spent twenty minutes without realizing it. She reflected: *"One need this time, really [...]. Because it was very loud from the very beginning. At the time when I put the headphones on, there is already sound and stuff, and as I must adapt to the system, it is as if I was already late, I needed to understand it quickly"*. Another one, who had experience in drumming, inquired his own sense of control with the sticks: *"It took me time, I took actions, but when I stopped, I've started to question: Was it just an illusion?"*, he said.

Adapting one's motion Beyond reflecting on the system, the sticks had the four participants reflect on their own movement practice. The grasping affordances of the sticks had two participants feel an assymetry in their motion: *"It was really interesting to move with the sticks. [...] I rapidly realized that I was very lateralized, in the sense that my right hand dominates, and I am not trained in drumming"*, one participant analysed. A third participant also modified its movement exploration, this time in relation to the act of drumming: *"As soon as I saw it, I thought that I can play [the drums]. Thus, I begin to sit down"*, she commented.

Appropriating the sonic mapping Three of the four participants appropriated interaction with the *somasticks* in relation to the sounds that they produced. *"For me, there is something very attractive in the fact that there are moments of synchrony and moments of autonomy. And in its autonomy, there were lots of variation still, so it was always interesting to listen to, and move along with"*, one expert reported about the interactive mapping. Another one seemed to understand the windowed implementation of the unsupervised learning mapping: *"I understood that there was a reaction at this level. [...] I understood that there was a delay. This meant that if I am doing the same thing during three seconds, the totality [sic] will come after"*, she said. The

last one of the three was first *"in quest of silence"* trying not to move, then seemed to gain control over the wavelet-based implementation of the sticks: *"There is a kind of obstacle course that appeared. [...] It was when I made small movements that I managed to refocus on what I was doing, and to take back control over the system"*, she said.

Link with Expertise Despite the playful moment of interacting with the sticks, the two remaining participants did not manage to understand the sticks' functioning. One of them related this by his lack of expertise in drumming: *"I do not know all these sensations [of drumming], so it was a bit harder for me to find the good movement qualities"*, he analysed. The other was rather disturbed by the physicality of the sticks, which had him want to play them on an actual drum: *"I think it could be great, even to refine the system in a kind of design process, to play it with something, you know, just, like an augmented instrument"*, he said.

3.4 Discussion

In this section, we discuss challenges related to machine learning implementation, as well as insight for general human-machine interaction.

3.4.1 Challenges in Implementation

Our current system relies on one learning model (GMM) and a slider-based GUI. Future work may implement an online expectation-maximization algorithm for continuous, optimized learning and inferring, and investigate interactive visualizations of Gaussian distributions to let users interact in an audiovisual augmented reality setup. Information-theoretic measures such as entropy may be better formalized and implemented to enable alternative mapping explorations.

Also and importantly, we would like to prototype other unsupervised learning models, allowing for even more diverse musical uses. For example, a current limitation of the *Gaussian Mixture Model* is that it considers each new input as independent from previously-observed data points. Such a property might not be suitable to human movement, as dynamics are deemed of prime importance when dealing with qualities of corporeal expressiveness [Leman, 2016]. Therefore, modelling dynamic patterns in gestural data could be a promising approach for generating sequential musical output that would be stylistically coherent with users' bodily expression. Future work may study adaptive dynamical systems to both model user-specific movement qualities and to generate continuous navigation trajectories [Kulic et al., 2007]. Another approach would be to study a reactive factor oracle [Chemla et al., 2015] to let users either shape a training set of movement patterns, or guide a discrete navigation through this training set.

Finally, our current implementation does not provide humans with a completely continuous way to interact with machine learning. If the number of GUI actions has been reduced from previous interactive supervised learning systems, users still have to specify whether they would like to record, delete, or window stream data during their performance. Other memory processes may be investigated to allow automatic

recognition of physical actions taken by users [Kulic et al., 2007], thus mediating embodied musical interactions more fluidly.

3.4.2 Designing for Human Exploration

Observations harvested during the demonstration of the *somasticks* suggest that the design through performance workflow importantly rely on listening abilities in relation to motion. Such an online listening workflow helped expert performers explore a variety of musical situations, which inspired their creative practice. As such, design through performance may be analysed as a workflow aimed at improving flow in creative practice [Csikszentmihalyi, 1997]. Future work may investigate interaction over longer periods of time to study how *somasticks* support flow states in relation to appropriation and learning time.

While the *somasticks* only relied on the “guiding" interaction mode to provide such flow, the “shaping" mode remains implemented under the same experiential workflow (action-perception loops emerging from physical interaction with sound. This could enable alternative exploratory uses by performers. For example, one could add data to the training set indefinitely to create a mapping that would progressively “freeze" once having recorded enough data. Also and perhaps surprisingly, the “Delete" action actually produces sound: one could imagine a performance where “Delete" gestures would act as control mechanisms for sonic events. Several new interaction styles could thus be explored with our tool, each of them placing corporeal engagement with sound as the main point of focus.

In the context of this thesis, we decided not to go deeper in the study of motion-sound mapping to reflect on the broader idea of supporting flow with machine learning. Specifically, we became interested in the concept of supporting *human exploration*. This new strong concept [Höök and Löwgren, 2012] emerged from the design of our model prototype aimed at movement-based human performance. As a matter of fact, the online adaptation of our model prototype enabled human exploration of various motion-sound mappings, as well as of various movement strategies in relation to sound. We were interested in researching other applications of human-centred machine learning that could address the issue of human exploration for other musical tasks. This will be the theme of the two next chapters.

4

Designing With Reinforcement Learning For Sonic Exploration

This chapter presents the second study led in the thesis. Based on the insight harvested in our study of motion-sound mapping, we decided to focus on the musical task of sonic exploration. Sonic exploration consists in listening to many different sound designs to converge to a final sound design. Our idea is to use machine learning to enable autonomous sonic exploration in close partnership with humans.

Sonic exploration is a multimodal task that involves both sound listening and interface actioning from humans. We decided to address these two modalities through two separate studies. In this chapter, we adopted an experimental method [Lazar et al., 2017] to study how humans perceive autonomous machine exploration through sound listening only. Figure 4.1 schematizes the three steps of our study. The issue of interface actioning in sonic exploration will be addressed in Chapter 5.

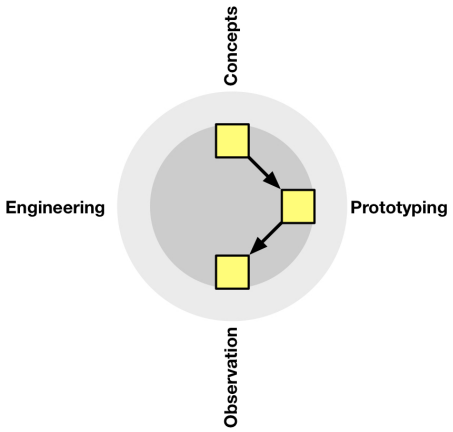


Figure 4.1: The experimental method for our second study.

Section 4.1 describes **concepts** on sonic exploration, covering musical context along with our proposed machine learning technique. Section 4.2 describes the model **prototyping** step of our design process. Section 4.3 reports on **observations** led in a controlled setup, studying human perception of our model in the task of sonic exploration.

This work was published and presented as paper and presentation at the *15th Sound and Music Computing Conference (SMC 2018)* in Limassol, Cyprus [Scurto et al., 2018a]. Concepts were published and presented as paper and presentation at the *Journées d'Informatique Musicale (JIM 2018)* in Amiens, France [Scurto and Bevilacqua, 2018].

4.1 Concepts

The introduction of the computer in music has brought a wealth of novel practices around sound and music. Cutting-edge technologies have been developed for sound synthesis, processing, analysis, and control, enabling the emergence of new music works, practices, notations, and performances.

However, such new technologies remain hard for people to appropriate. While music computing is now taught in many music institutions, it still suffers from its apparent complexity. As a consequence, many musicians stay attached to their classical practices and resign to appropriate music computing, while many musicians-to-be never take a chance to explore music computing systems. This is paradoxical as computing has become ubiquitous in the last ten years. Can we think of a computer tool that would facilitate exploration and appropriation of music computing systems? We focus on a particular use case of music computing appropriation. The use case refers to the situation where a user makes use of a music computing system on his or her own, outside educational institutions—*e.g.,* amateur music practices led in a home studio. In this situation, appropriation issues arise during the first uses of the music computing system.

Appropriation in Music Computing

Starting using a system generally requires gathering information on it—in a passive learning setup. Information on a system can be found in various media, from the most straightforward (*e.g.*, a text or video tutorial) to the most technical (*e.g.*, a research paper), as well as through online discussion (*e.g.*, an Internet forum). These activities are often time-consuming: users first have to find relevant information (which can be hard for obsolescent systems), then to filter it (*i.e.*, find what is useful for a specific goal), supposed that they have a specific goal in mind. Overall, time spent on passively learning hinders users to interact with sound and music directly, which might drive them away from experimenting with the system.

Another option, which could be combined with the first, consists in starting interacting with the system's interface from scratch—in an active learning setup. Interacting with the interface implies trying many different actions directly to understand the functioning of the system (in a trial-and-error fashion), and eventually to achieve

a specific goal. In our case of music computing, these activities are crucial as it is important for users to actively control sound so as to strengthen action-perception loops [Leman, 2008]. Yet, it is also possible that users get discouraged in interacting with the interface if they get too much error during their trials. Notions of appropriation [Zappi and McPherson, 2014] and novice to expert transition [Cockburn et al., 2015] are thus crucial in the design of interactions at stake in a given interface.

Human-Computer Interaction in Music Computing Systems

We identify two drawbacks of current music computing systems in the context of appropriation. First, most music computing systems' interfaces can look quite intimidating for completely novice users. Some of them directly derive from their analog ancestors (*e.g.*, sound synthesis engines, see Figure 4.2, or spatialization tools [Carpentier et al., 2015]): they are thus designed for expert users, not to facilitate interaction for novice users. There do have been attempts to improve interface accessibility of these systems [Cartwright et al., 2014, Cartwright and Pardo, 2014, Schwarz et al., 2006]. Yet, these interface simplifications are often done to the detriment of the system's abilities: after having appropriated some tasks, humans become limited by the interface's lack of sophistication.



Figure 4.2: Example user interface of a music computing system (here, the u-he Bazille VST[1]).

The second drawback is that music computing systems often rely on complex models and representations that are not directly linked to sound. For example, improvisational systems [Assayag et al., 2006] require humans to have knowledge on mathematical models, as well as on musical representations, to understand how parameters relate to sound. Similarly, musical environemnts such as Max or PureData

[1] http://www.u-he.com/cms/bazille

require to learn new representations as well as programming to start interacting. This drawback also concerns expert musicians—who might know a certain amount of musical parameters, but do not know how they relate to new mathematical parameters—, and constitutes a major issue for music pedagogy—as such music computing systems have been shown promising for instrument learning [Cont, 2008] and embodied practice [Bevilacqua et al., 2007].

4.1.1 Musical Task: Sonic Exploration

We believe that designing for exploration could be a way to improve appropriation of these interfaces and representations. Exploration is the early phase of learning during which a human iteratively acts on an interface and receives feedback information, allowing him or her to gradually grasp the system's functioning and qualities. As discussed previously, it is a crucial phase in appropriation regarding learning and skill development, as good or bad initial experience will determine the future degree of motivation and involvement of a learner term for a given task [Cockburn et al., 2015]. By supporting exploration phase, our wish is to lower the threshold for learners to directly interact with the system and sense its abilities, paving the way for further understandings of how the system actually works.

Situating the Musical Task

Exploration of music computing systems is a multimodal task. One may distinguish between sonic exploration and interface exploration. Sonic exploration consists in listening to several features of a sound to learn *what* qualities the created sound should possess. Interface exploration consists in acting on the parameters of an interface to learn *how* to create this sound. While the two modalities are intertwined, we decided to first focus on the sonic exploration task only[1]. Our wish was to test whether listening could be used by non-musicians as an expressive modality to explore and appropriate musical interfaces [Leman, 2016].

As a remark, sonic exploration tasks also concern a wide variety of expert musicians. For example, composers explore various sketches of their musical ideas to write a final score; performers explore different playing modes to shape an instrument's tone; sound designers explore several digital audio parameters to create unheard-of sounds [Monache et al., 2010, Delle Monache et al., 2018]. Our interest lies in designing techniques that may support sonic exploration for both musicians and non-musicians. Specifically, we are interested in developing interaction modalities that may allow humans to leverage their listening perceptive abilities to explore sounds.

Machine-Centred Approaches

In the last decade, the field of Interactive Machine Learning has enabled exploration of music computing systems by emphasizing embodied interaction with sound. Supervised learning is the technique that has been mainly investigated in interactive machine learning, especially for motion-sound mapping [Françoise et al., 2014, Carami-

[1] We address the musical task of interface exploration in Chapter 5.

aux and Tanaka, 2013]. Human interaction with supervised learning consists in demonstrating example gestures to the learning algorithm so that it can learn to recognize them on the fly. Exploration within supervised learning either takes place during the training phase (where humans can experiment with several examples of different gestures), or during the running phase (where they can explore interpolations between given examples) [Fiebrink et al., 2010]. This two-phase workflow has been shown useful for a number of tasks; however, it has been shown limiting in some cases, for example when humans want to slightly modify a given design [Scurto et al., 2016].

Sequential adaptation have been investigated for interacting with autonomous agents [Assayag, 2014], in a design through performance fashion [Scurto et al., 2017]. Human interaction with autonomous agents consists in generating example musical content for guiding agents' musical behavior. Exploration within autonomous agents mainly consists in continuous musical improvization with the agent [Nika et al., 2017]. This online learning workflow has been shown useful for performance cases (which require continuous generation and reactivity) [Assayag et al., 2006, Pachet, 2003] but may not be fully adapted to more general, "offline" design cases. Crucially, it still coerces humans into demonstrating expert musical examples to explore new behaviors, which might prevent non-musicians from interacting with these systems.

Human-Centred Approaches

From a pedagogical point of view, the potential of interactive machine learning systems has been identified, yet little exploited. For gestural control of sound, they have been cited as allowing "learners to experience components of higher-level creativity and social interaction even before developing the prerequisite sensorimotor skills or academic knowledge" [Morris and Fiebrink, 2013]. Interestingly, novel application domains, such as music therapy and musical expression for people with disabilities, have also emerged [Scurto et al., 2016]. We believe extending interactive machine learning approaches to other music computing systems could constitute an opportunity to widen the reach of more music computing practices to more people.

Many expert musicians reported that interactive supervised learning systems offered space for exploration [Fiebrink et al., 2010], often personifying them as *collaborative partners* because of their ability to learn implicit musical properties similarly to a human musical collaborator [Fiebrink and Caramiaux, 2016, Pachet, 2003]. Recently, new interactions with supervised learning algorithms have been investigated to support exploration for both novice and expert users. Scurto et al. implemented a machine learning tool able to generate many alternative user-adapted mappings from only one motion stream [Scurto et al., 2016]. This workflow avoided users to reflect on what examples they should demonstrate for reaching a goal: rather, it enabled them to focus only on subjective, evaluative exploration of many prototypes. Users valued the space for exploration offered by such autonomous generation abilities. However, they expressed a lack of control over the system, as generation remained fully autonomous and not adaptive—neither sequentially, nor subjectively.

4.1.2 Proposed Technique: Reinforcement Learning

To assist sonic exploration, we find it relevant to use the metaphor of transmission of knowledge between humans. Consider a human that has an idea but does not know how to convert it in a concrete realization. Usually, the human will ask assistance to a second human to realize this conversion—we call it the assistant. Iterative interaction between the two humans takes place, during which the assistant takes actions on the system and the human gives feedback on it—until converging to a final design. Our idea is to have a machine learning agent take the role of the assistant: the agent acts on the system, upon which the human gives feedback. The agent thus explores design possibilities in collaboration with the human, letting the human focus solely on aligning their conceptual space with the perceptual space offered by the agent—postponing the sensorimotor and/or academic learning phase to a later phase.

Reinforcement Learning

We propose to investigate the interactive use of a specific category of machine learning, called reinforcement learning. Reinforcement learning defines a formal framework for the interaction between a learning agent and an environment in terms of states, actions, and rewards [Sutton and Barto, 2011]. At time t, an agent senses its environment through an observation called state S_t (typically, a vector of discrete parameters), and on that basis takes an action A_t on it (typically, a set of discrete modifications on these parameters). At time $t+1$, in response to its action, the agent receives a reward R_{t+1} from the environment, as well as a new state S_{t+1}. From this information, the agent iterates interaction, progressively learning how to optimize interaction with the environment so as to maximize the total amount of reward it receives over the long run.

Reinforcement learning differ from supervised (and unsupervised) learning categories. For the latters, learning typically occurs offline on the basis of a static training dataset, which is a set of labeled (or unlabeled) examples we would like the system to generalize behaviour from. In reinforcement learning, the agent learns online by directly interacting with its environment. As a result, a reinforcement learning agent must always balance between exploration and exploitation to improve its learning – exploration meaning trying new actions to discover which ones yield the most reward, and exploitation meaning choosing the best actions in terms of reward at the time of computation.

Proposed Workflow

We are interested in *interactive* approaches to reinforcement learning, such as shown in Figure 4.3. In this approach, the reinforcement learning agent receives the reward signal directly from the human feedback. Research in other fields such as robotics [Thomaz and Breazeal, 2008] and computer science [Knox and Stone, 2009, Christiano et al., 2017] have made huge progress toward the development of interactive agents capable of learning specific tasks from human feedback. These agents support sequential adaptation without needing example demonstrations, but only by receiving human feedback as subjective evaluations of the autonomously-generated behaviors.

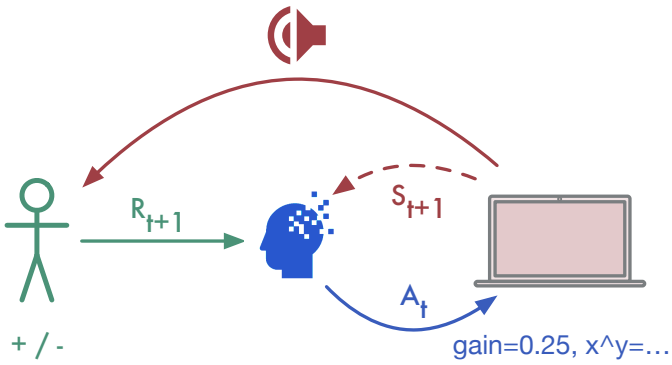


Figure 4.3: Graphical outline of an *Interactive Reinforcement Learning* framework. Interaction between the human and the music computing system is mediated by a learning agent (in blue). At time t, the agent acts directly on the system's parameters. At time $t+1$, the system generates a new state (for example, a sound), that is subjectively evaluated by the human through feedback. By iterating the loop, the agent learns how to explore.

There is still relatively few works investigating reinforcement learning in the field of music computing [Fiebrink and Caramiaux, 2016]. Derbinsky et al. [Derbinsky and Essl, 2012] proposed to apply reinforcement learning to rhythm modelling, in a goal-oriented task. Our approach differs from these works, in the sense that we focus on interactive agent teaching from the human point of view (how it is "efficient" for the human, not necessarily for the agent). Also, it extends research in other fields on interactive reinforcement learning in the sense that the open-ended task of sonic exploration may have different properties than the goal-oriented tasks usually investigated—such as learning how to play Tetris [Knox and Stone, 2009].

We propose to conceptualize the application of interactive reinforcement learning to human exploration as "co-exploration" (see Figure 4.4). We are inspired by prior work studying phenomenons of co-adaptation between humans and machines, which were shown to reduce the human's cognitive overload during the carrying of a task [Mackay, 2000]. Here, co-exploration stands for collaborative human-agent exploration of a given music computing system. The human explores the sonic possibilities offered by the system (progressively learning aesthetic qualities in the system), while the agent explores the computational parameterizations of the system (progressively learning which system parameters are relevant for the human). Design through co-exploration encapsulates the possibility for a (possibly novice) human to create a musical artifact from a (possibly unknown) computer system by collaborating with a learning agent in the exploration of design possibilities. A typical scenario would imply the agent generate an initial random solution to the human, who would progressively shape it through her preferences. As we saw it, human interaction with the music computing system is mediated by an agent. Concretely, this means that the human does not interact with the system's interface: an agent does it instead. Instead, the human focuses on giving evaluative feedback on the agent's actions, judging the

system's output on a perceptual level. Potentially, this feedback could be of any type—be it text-based, demonstration-based, or physiological.

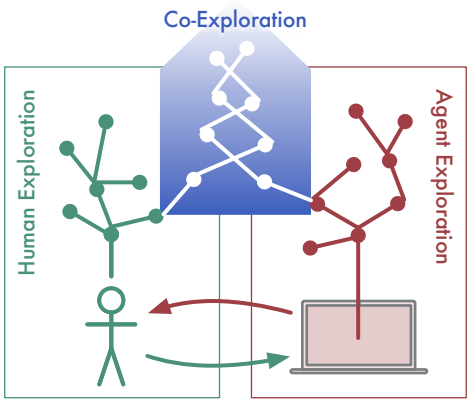


Figure 4.4: The *Co-Exploration* workflow. In a standard situation (behind), a human explores a system by iteratively acting on it. In co-exploration (front), a learning agent explores a system in parallel to the human.

4.2 Prototyping

In this section, we describe the model prototype that we implemented to apply interactive reinforcement learning to the musical task of sonic exploration.

4.2.1 Human Data: Feedback

Our initial wish is to allow humans to explore sound by only communicating positive or negative feedback data. Advantages are threefold. First, it could encapsulate several kinds of feedback in one unique format, such as general advice (*e.g.* "this is good", "this is bad"), implicit knowledge (*e.g.* "do it more like this", "don't go that way"), as well as explicit specification (*e.g.* "this is exactly what I want", "never show me this again"). Second, it could be expressed relatively easily (compared to text-based feedback, which forces users to create a concrete verbalization of what they want). Third, it could give a sense of agency to the human (compared to physiological feedback, which most humans do not control). Overall, we hypothesize that communicating such high-level feedback could facilitate musical exploration of a system compared to specifying its low-level parameters.

To implement interaction with such agents, we must modify the formal framework defined above. We propose, along with previous works [Thomaz and Breazeal, 2008, Knox and Stone, 2009, Christiano et al., 2017], that a human would be responsible for giving reward to the agent (see Figure 4.3). Our hypotheses are that the numerical reward may constitute a feedback channel from the human to the agent (respectively giving positive, zero, or negative reward for positive, neutral, or negative feedback),

and that interactively communicating feedback toward the environment following the agent's exploration path may support human exploration.

4.2.2 Algorithm: *Sarsa*

We implemented *Sarsa*, which is a standard algorithm to learn how to act in many different environment states [Sutton and Barto, 2011]. It differs from multi-armed bandits, which learns how to act in one unique environment state [Lomas et al., 2016]. Formally, the environment is constituted by the parameters of some sound synthesis engine, and the agent iteratively acts on them. Computationally, we considered the state space $\mathcal{S} = \{S\}$ constituted by all possible parameter configurations $S = (s_1, ..., s_n)$, with n being the number of parameters, and $s_i \in [s_{min}, s_{max}]$ being the value of the i^{th} parameter living in some bounded numerical range (for example, s_i can control the level of noise normalized between 0 and 1). We defined the corresponding action space $\mathcal{A}(S) = \{A\}$ as moving up or down one of the n parameters by one step a_i, except when the selected parameter equals one boundary value:

$$A(S) = \begin{cases} \pm a_i & \text{for } s_i \in]s_{min}, s_{max}[\\ +a_i & \text{for } s_i = s_{min} \\ -a_i & \text{for } s_i = s_{max} \end{cases} \tag{4.1}$$

An ε-greedy method defines the autonomous exploration behaviour policy of the agent—how it may act by exploiting its accumulated feedback while still exploring new unvisited states [Sutton and Barto, 2011]. It consists in having the agent take an optimal action with probability ε, and reciprocally, take a random action with probability $1 - \varepsilon$. For example, $\varepsilon = 1$ would configure an always exploiting agent—*i.e.,* always taking the best actions based on accumulated feedback—, while $\varepsilon = 0$ would configure an always exploring agent—*i.e.,* never taking into account the received feedback.

4.3 Observation

As a first step toward co-exploration, we led a controlled experiment with human participants. Our aims were to: Test how humans perceive interaction with a reinforcement learning agent for sonic exploration; Understand how reinforcement learning elements may influence humans leading sonic exploration.

4.3.1 Setup: Controlled Experiment

Participants

We recruited 12 participants (average of 26.9 years old, $\sigma = 7.44$, 5 Female and 7 Male). Half of them were music computing practitioners. The other half did not have training in music computing, nor other sound-related creative practice. All of them reported normal hearing.

Task

The basic task of the study was to guide an agent through a sound space, from the lowest to the brightest sound. At each step of the task, the agent would generate a new sound. If the new sound was brighter than the previously generated one, participants had to give positive feedback to the agent. In any other cases (lower or similar brightness), participants had to give negative feedback to the agent. The task automatically ended in two cases: either the brightest sound was reached, or it was not reached after a maximum number of steps (we set it to 150).

At the end of the task, participants were asked to rate their perception of the agent according to three aspects related to collaboration. The first aspect was the degree of agency provided by the agent through feedback (*"did the agent seem to take into account your feedback in a reactive manner, or did it seem to act completely independently?"*). The second aspect was the degree of assistance provided by the agent throughout the task (*"did the agent seem to generate sounds that were brighter, or did it seem not to be of any help in going in this direction?"*). The third and last aspect was the degree of easiness of the task (*"overall, did the task seem to be very easy, or very difficult?"*).

Agents

Three types of agents were evaluated: "random", "balance", and "exploit". These correspond to three different degrees of exploration ($\varepsilon = 0$: the agent only takes random actions; $\varepsilon = 0.5$: the agent balances random action selection with feedback-based best action selection with probability 0.5; $\varepsilon = 1$: the agent only selects the best actions as indicated by user feedback). Other agent parameters were fixed so that exploration would be the sole varying factor.

Musical Environments

Sounds were generated through a FM synthesis[2] engine (implemented in Max/MSP), with two discretized parameters. The first parameter, called modulation index, could take ten values ranging from 3 to 70; the second parameter, called harmonicity ratio, could take three values ranging from 0.98 to 1.02. The resulting sound space thus had 30 possible states, corresponding to 30 static sounds. As previously explained, the agent's possible actions consist in moving up or down one of the two parameters. For the sake of the experiment, we normalized sound loudness empirically so they perceptually appear of equal intensity, and we set sound duration to 500 ms.

Based on this sound space, we designed two environment models in close relationship with the task's goal: "unobstrucked", and "obstrucked" (see Figure 4.5). In the unobstrucked environment, brightness increases linearly with modulation index: highest brightness thus corresponds to highest index value. We expect "balance" and "exploit" agents to be more collaborative than "random" agents through their ability to learn and select the best actions.

[2] *Frequency Modulation synthesis* (a classic algorithmic method for sound synthesis [Chowning, 1973]).

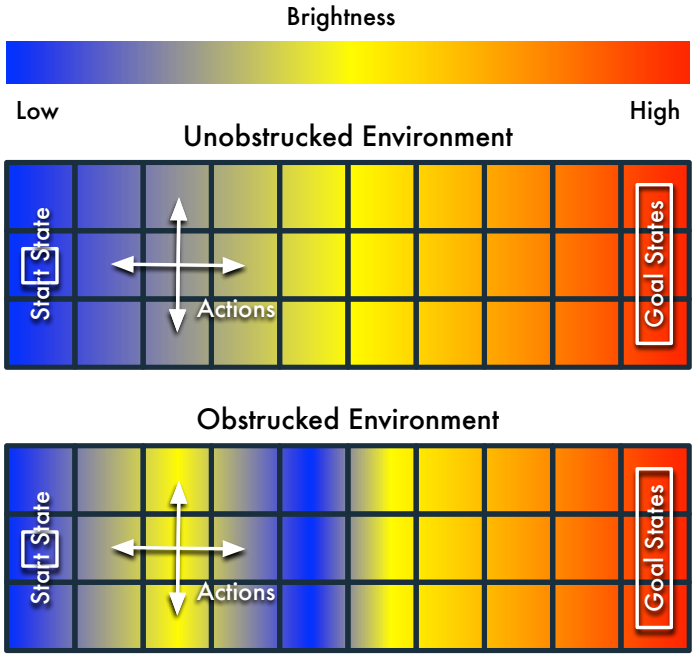


Figure 4.5: The two environment models designed for our experiment. Top: Unobstrucked environment, where brightness varies linearly. Bottom: Obstrucked environment, where brightness varies nonlinearly.

In the obstrucked environment, brightness varies nonlinearly with modulation index: highest brightness still corresponds to highest index value, but a local maximum lives at one third of the scale. Our hypothesis is that "exploit" agents would remain stuck in this local maximum, whereas "balance" agents would overcome it through their ability to explore. We thus expect "balance" agents to be more collaborative than "random" and "exploit" agents.

Procedure

The experimental session consisted of a familiarization phase and an experimental phase.

Participants first had to read the task's instruction and could ask the experimenter for clarification if necessary. Then, they had two test tasks in the unobstrucked environment with two types of agents (one "exploit", then one "random") to familiarize with the range of sounds and agent behaviors at stake. Sounds were presented as pairs to participants (using headphones), so as to facilitate brightness comparison between the previously-generated sound and the new one. Participants could listen to a pair of sounds as many time as they wanted to (using a keyboard key) before giving positive or negative feedback to the agent (using left or right arrow keys). Once a task was over, participants had to rate the agent's behavior for each of the three previously-described aspects on 9-point Likert scales (using the mouse and interactive sliders).

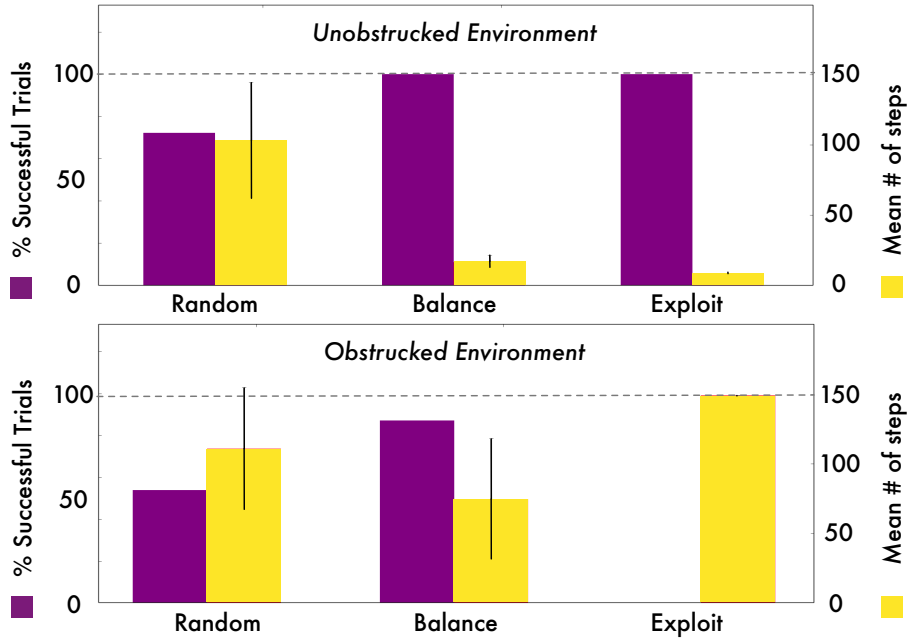


Figure 4.6: Results for synthetic trial data.

We asked participants to use the full scales as much as they could.

Once this phase was over, participants could start the experimental phase. The first stage only concerned the unobstrucked environment: participants were asked to guide and evaluate each of the three types of agents within it. For improving consistency, participants made three trials with each of the three agents. A stage thus consisted in nine tasks that were randomized in order. Finally, the second stage only concerned the obstrucked environment: similarly, participants guided and evaluated the three types of agents three times each, in a random order. Participants were allowed to take a break at any time during the session, which lasted one hour on average.

4.3.2 Results: Statistical Analysis

For each participant, we recorded step-by-step data (time, states, actions, feedback and ratings), as well as audiovisual data of users. Prior to analysing them, we report on synthetic data generated before the actual experiment.

Synthetic Trial Data

We programmed synthetic feedback users of same number as participants to generate a benchmark on how agents should ideally behave in our two environment models. This case corresponds to participants giving perfectly consistent feedback.

We measured the percentage of successful trials (which reflects the probability of reaching the goal), as well as the mean number of steps taken in a trial (which reflects a trial's duration), for each type of agent and in each of the two environ-

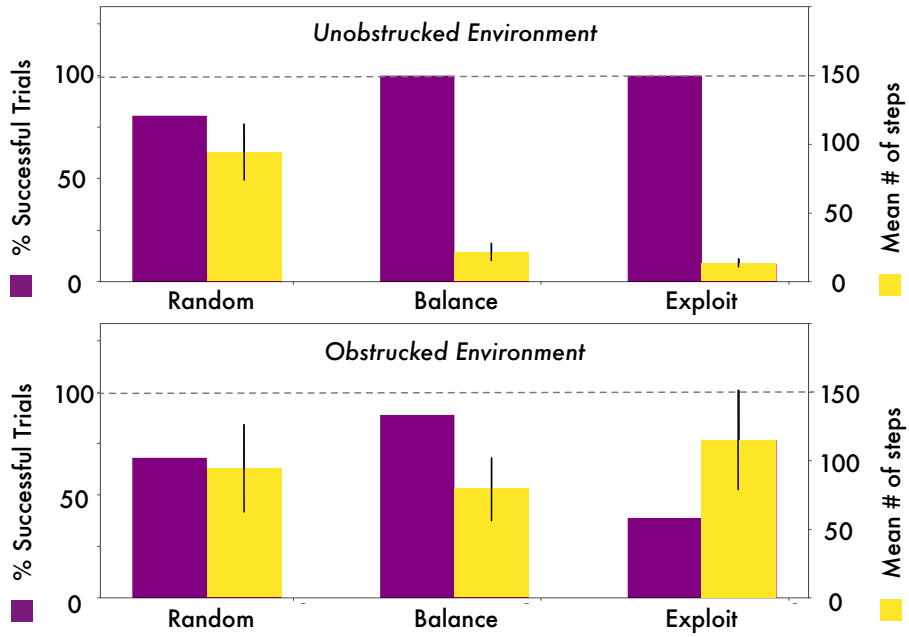


Figure 4.7: Results for participants' trial data.

ments (see Figure 4.6). For each environment, we submitted each measure to a one-way ANOVA with agent exploration as the within-subject factor. In the unobstrucked environment, the effect of exploration was significant for both percentage of successful trials [$F(2, 22) = 8.83, p < 0.001$] and mean number of steps [$F(2, 22) = 91.3, p < 0.001$]. Planned contrasts showed that both measures significantly differed for "balance" and "exploit" agents compared to "random" agents.

Likewise, in the obstrucked environment, the effect of exploration was significant for number of successful trials [$F(2, 22) = 44.7, p < 0.001$] and mean number of steps [$F(2, 22) = 26.3, p < 0.001$]. Planned contrasts showed that both measures significantly differed for "balance" agents compared to "random" and "exploit" agents.

Participants' Trial Data

We first measured participants' feedback behavior. In the unobstrucked environment, participants gave a mean of 393 feedback every 1.91 s, with 96.3% being correct. In the obstrucked environment, participants gave a mean of 879 feedback every 1.84 s, with 98.0% being correct.

Similarly to synthetic users, we measured the percentage of successful trials, as well as the mean number of steps taken by each of the three agent types, in each of the two environments (see Figure 4.7). We used the mean of all trials in each condition for each participant. For both environments, we submitted both measures to a one-way ANOVA with agent exploration as the within-subject factor. In the unobstrucked environment, the effect of exploration was significant for percentage of successful trials [$F(2, 22) = 6.49, p < 0.005$] and mean number of steps

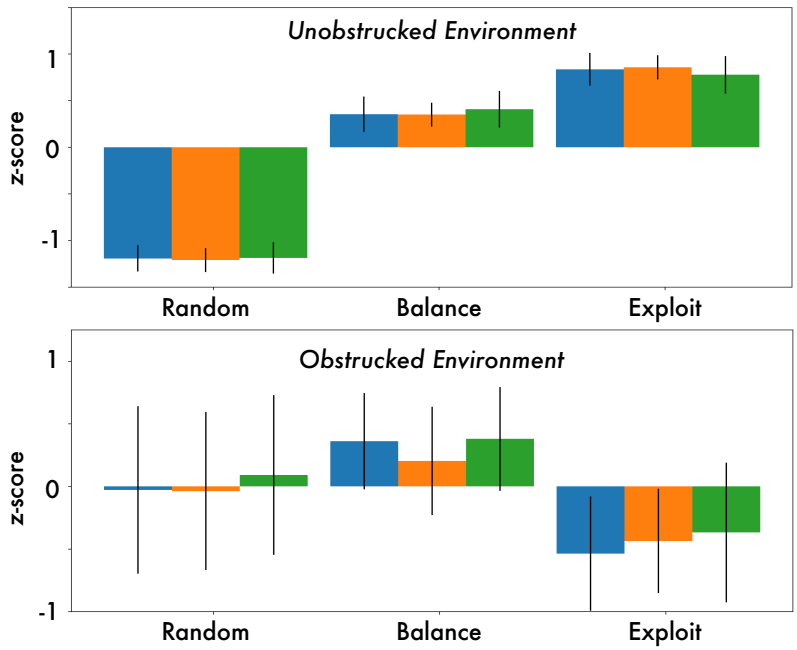


Figure 4.8: Results for participants' evaluation data. In blue: agency. In orange: assistance. In green: easiness.

$[F(2, 22) = 130.3, p < 0.001]$. Planned contrasts showed that both measures significantly differed for "balance" and "exploit" agents compared to "random" agents.

Likewise, in the obstrucked environment, the effect of exploration was significant for percentage of successful trials $[F(2, 22) = 8.16, p < 0.002]$ and mean number of steps $[F(2, 22) = 3.62, p < 0.03]$. Planned contrasts showed that both measures significantly differed for "balance" agents compared to "random" and "exploit" agents.

Participants' Evaluation Data

We computed the standard score (also called z-score) for each evaluation ratings in each environment to compare participants on the same scale (see Figure 4.8).

For each environment, we submitted each z-score to a one-way ANOVA with agent exploration as the within-subject factor. In the unobstrucked environment, the effect of exploration was significant for all three perceptual aspects ($[F(2, 22) = 429.3, p < 0.001]$ for agency; $[F(2, 22) = 767.3, p < 0.001]$ for assistance; and $[F(2, 22) = 335.2, p < 0.001]$ for easiness). Planned contrasts showed that all three perceptual ratings were significantly higher for "balance" and "exploit" agents than for "random" agents.

Likewise, in the obstrucked environment, the effect of exploration was significant for for all three perceptual aspects ($[F(2, 22) = 8.32, p < 0.002]$ for agency; $[F(2, 22) = 4.53, p < 0.02]$ for assistance; and $[F(2, 22) = 5.26, p < 0.02]$ for easiness). Planned contrasts showed that all three perceptual ratings were significantly higher for "balance" agents than for "random" and "exploit" agents.

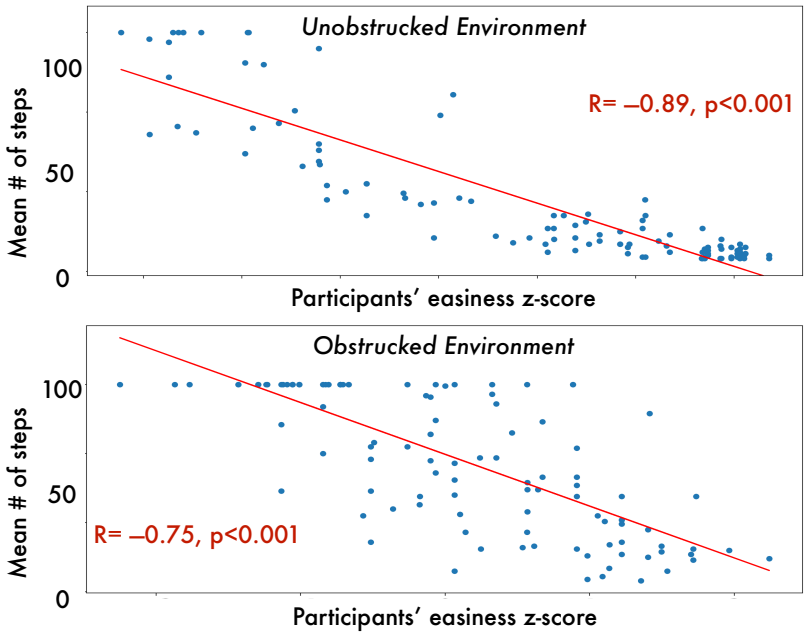


Figure 4.9: Results for participants' ratings versus task parameters.

Finally and as shown in Figure 4.9, we measured that participants' perception of task easiness was correlated with the total number of steps taken by all types of agents, in both environments.

4.4 Discussion

In this section, we discuss our experiment's results and extract implications for future investigations of interactive reinforcement learning applied to sonic exploration.

4.4.1 The Usefulness of Balancing Exploitation with Exploration

Synthetic trial data

We first look at synthetic trial data to analyse agents' ability to reach a goal in a non-interactive setup. In the unobstrucked environment, as expected, all agents that took into account feedback ("balance" and "exploit") always succeeded in reaching the goal, with "exploit" agents being the fastests as they took the best action at each step; "random" agents reported the worst performance, succeeding only two thirds of the time with lower speed. In the obstrucked environment, conversely, "exploit" agents never succeeded in reaching the goal. As expected, they remained stuck in the local maximum that we designed. In this case where an obstacle blocks the way to the goal, "balance" agents remarkably outperformed other agents in both speed and number of

success. This proves that agents' balance between exploitation and exploration may be useful for reaching a goal in environments of varying complexities.

Participants' trial data

Participants' trial data differ from synthetic trial data because of imperfect feedback occasionally given by users. Despite this difference, agents took exploration paths that were similar to those generated with synthetic users in five out of six agent-environment cases, as shown in Figure 4.7. In the remaining case of "exploit" agents exploring the obstrucked environment, one third of the trials were successful, which means that agents unexpectedly managed to overcome the obstacle that we designed to reach the goal. This proves that agents can take different paths in an interactive setup where users make feedback mistakes.

4.4.2 The Influence of Exploration Path on User Perception

Perceiving collaboration

We now analyse participants' subjective evaluations to better understand how exploration might be perceived by users. First, we observe that participants' ratings had more variability in the obstrucked environment than in the unobstrucked environment. This suggests that an environment's complexity may strongly influence how humans perceive agent exploration. Second, we noticed that participants rated down "exploit" agents in the obstrucked environment, even if one third of them succeeded in reaching the goal, as we previously discussed. This proves that the path taken by agents during exploration may be more critical to how collaborative agents are perceived by users than the actual fact of reaching the goal.

Looking more in detail to participants' ratings, we can see that "balance" agents were the only type of agents that were perceived as being the most assistive in both environments, thus reflecting their quantitative usefulness. As expected, "random" agents were perceived as providing the less agency in the unobstrucked environment: this suggests that participants may be able to perceive when an agent learns along its path—in other words, there was no "placebo effect" toward agents' artificial intelligence. Finally, even if "exploit" agents formally take the best action at every step as defined by participants' feedback, this may not be perceived by participants, as their ratings of agency shows (see Figure 4.8, bottom). This confirms that an agent's internal functioning may not be properly perceived by humans, whose perception might be more influenced by the path taken by agents in a given environment. Results shown on Figure 4.9 seem to confirm this statement, as one of the evaluation ratings correlates with one of the task parameters, regardless of the type of agent at stake.

Personifying agents

Interestingly, audiovisual recordings show that all participants personified agents depending on their perceived collaboration. For example, agents that took relatively direct paths to the goal provoked positive reactions (such as *"it understood right away"*)

and adjectives (e.g., *"nice"*, or *"careful"*). On the other hand, agents that took more complex paths—such as "random" agents, or "exploit" agents that remained stuck in the obstacle—inherited depreciative reactions (e.g., *"it doesn't listen to me"*, or *"it seems light-headed"*) and adjectives (e.g., *"idiot"*). This might be a first clue—to some extent—for stating that feedback-based interaction may encourage users to perceive agents as embodied partners—in some cases able to act as collaborators.

4.4.3 Towards Co-Exploration

The issue of human moving goals

In our experiment, we forced participants to follow a fixed feedback strategy: this might limit the reach of our experiment's results. Indeed, such feedback constraint might not be realistic in real-world exploration, mainly for two points: (1) users might change their feedback strategy, and (2) their goals might evolve over time. These situations are typical of real-world scenarios, where users may push systems in limit conditions [Jorda, 2005], or may want to explore several alternative strategies [Fiebrink et al., 2010]. Investigating these points constitute next steps toward turning our interactive reinforcement learning system (where the goal to be learned was fixed) into a co-exploration system (where the goal to be learned might evolve as the human uses the system).

Improving algorithms or interactions?

We identify two main directions for addressing these points—stressing that these directions should be complementary. The first option corresponds to investigate other reinforcement learning algorithms. As said, our current prototype implements the *Sarsa* algorithm, which is a standard method for reinforcement learning. Other approaches to learning may be better adapted to our co-exploration use case. For example, one may investigate methods that are robust to non-stationary feedback [Knox and Stone, 2009]. Alternatively, one may also investigate approximate policy learning algorithms [Sutton and Barto, 2011, Christiano et al., 2017] for learning relevant representations of an environment without having to explore it in its entirety.

The second option corresponds to design new interactions that may better fit interactive uses of reinforcement learning algorithms. As shown in our study, humans may not always perceive how a learning system internally works. In order to give more control to the human, one could imagine allowing humans to modify agent parameters during interaction, for example by actively choosing the degree of exploration they may need. Also, one could allow humans to go backwards in the agent's learning process, or to restart learning at any time, so as to give space for iterative, flexible exploration patterns [Resnick et al., 2005]. Again, all these developments are not contradictory, and we believe that both directions should be considered in future research.

Connecting agents to real-world systems and situations

Finally, our experiment focused on models of musical environments whose dimensionalities may not fully reflect those of standard music computing systems to be explored by users. Yet, we argue that investigating such models have provided useful insights on how agents would take exploration paths in real-world music systems.

In the context of this thesis, we decided to continue the application of reinforcement learning to sonic exploration with real-world VSTs, hoping to harvest complementary insights on our use case and pushing further the formalization of environments at stake in our co-exploration agents. Such studies might be an opportunity to investigate other qualitative methods for evaluating agents. Indeed, our experiment's results suggested that participants did not really differentiate each of the three perceptual aspects they had to rate, which in turn suggest that they may have a much global appreciation of how an agent interact with them. Borrowing approaches and methods from the field of Human-Computer Interaction (such as user-centered design through case studies and workshops) [Resnick et al., 2005, Fiebrink et al., 2011] might be essential for grasping such experiential aspects among humans and for leading such situated studies with agents.

5

Designing With Deep Reinforcement Learning For Synthesis Exploration

This chapter presents the third study led in the thesis. It follows on from our previous study in Chapter 4 by applying reinforcement learning to the musical task of "synthesis exploration". Synthesis exploration is the multimodal task of acting on a sound synthesis interface to lead sonic exploration.

We adopted a user-centered design method to iterate the design of reinforcement learning with humans [Beaudouin-Lafon and Mackay, 2009]. We decided to work with expert sound designers to get a specialized understanding of synthesis exploration parallel to design iterations. This enabled us to cover qualitative feedback in our model design process, which went through an engineering step to implement a deep reinforcement learning model (as shown in Fig. 5.1).

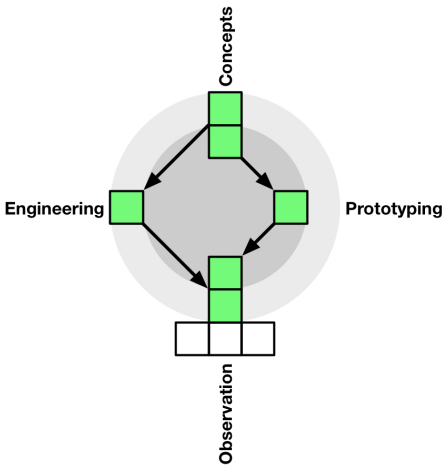


Figure 5.1: The user-centered design method for our third study.

Section 5.1 describes **concepts**, covering musical background on interface exploration, along with our proposed machine learning technique. Section 5.2 describes the first model **prototyping** iteration of our design process. Section 5.3 reports on **observations** led in a pilot study, respectively investigating expert users exploring VST interfaces and using our model prototype. Section 5.4 describes the second model **prototyping** iteration of our design process. Section 5.5 describes the model **engineering** led for our final software, called *Co-Explorer*. Section 5.6 reports on **observations** led in a creative workshop with our software.

Finally, we report on three **applications** of our model. Section 5.8 describes the application of the model to the design of a video game, called *Riding the Co-Explorers*, that we used as a demonstrator in a conference. Section 5.9 describes the application of the model to the design of a robotic object in a workshop, called *Behavioral Matter*. Finally, Section 5.10 describes the application of our model to the design of a computer music performance, called *ægo*.

This work was submitted as journal paper for *ACM Transactions on Computer-Human Interaction (TOCHI)* [Scurto et al., 2019b]. Engineering was done in the context of Bavo Van Kerrebroeck's Master's Thesis, supervised by Frédéric Bevilacqua and myself [Van Kerrebroeck, 2018]. The game was showcased as installation at the *19th International Society for Music Information Retrieval Conference* in Paris, France [Scurto et al., 2018b]. The robotic object was designed at the *Behavioral Matter* workshop in Paris, France. The computer music performance was published and presented as paper and piece at the *14th International Symposium on Computer Music Multidisciplinary Research (CMMR 2019)* in Marseille, France [Scurto and Chemla, 2019].

5.1 Concepts

Reinforcement learning defines a computational framework for the interaction between a learning agent and its environment [Niv, 2009]. The framework provides a basis for algorithms that learn an optimal behaviour in relation to the goal of a task [Sutton and Barto, 2011]. For example, reinforcement learning was recently used to learn to play the game of Go, simulating thousands of agent self-play games based on human expert games [Silver et al., 2016]. The algorithm, called *deep reinforcement learning*, leveraged advances in deep neural networks to tackle learning of a behaviour in high-dimensional spaces [Mnih et al., 2015]. The autonomous abilities of deep reinforcement learning agents let machine learning researchers foresee prominent applications in domains, such as transportation, healthcare, or finance [Li, 2018].

Yet, one important current challenge for real-world applications is the ability for reinforcement learning agents to learn from interaction with human users. The so-called *interactive reinforcement learning* framework has been shown to hold great potential to build autonomous systems that are centered on human users [Amershi et al., 2014], such as teachable and social robots [Thomaz and Breazeal, 2008], or assistive search engines [Athukorala et al., 2016a]. From a machine learning perspective, the main challenge lies in learning an optimal behaviour from small, non-stationary amounts of human data [Knox and Stone, 2009]. From a human-computer interaction perspective, an important challenge consists in supporting human appropriation

of algorithms' autonomous behaviours in relation to complex human tasks [Stumpf et al., 2009].

Our interest lies in investigating interactive reinforcement learning for human creative tasks, where a goal might not be well-defined by human users a priori [Resnick et al., 2005]. One such case of a human creative task is *exploration* [Hart et al., 2017]. Exploration consists in trying different solutions to address a problem, encouraging the co-evolution of the solution and the problem itself [Dorst and Cross, 2001]. For example, designers may produce several sketches of a product to ideate the features of its final design, or test several parameter combinations of a software tool to create alternative designs in the case where the product has a digital form. The creative, human-centred, use case of exploration fundamentally differs from standard, machine-centred, reinforcement learning use cases, where a problem is implicitly defined as a goal behaviour, before the agent actually learns to find a solution as optimal behaviour [Sutton and Barto, 2011]. It thus stands as an exemplary use case to study human interaction with reinforcement learning agents.

In this study, we aim at designing an interactive reinforcement learning system supporting human creative exploration. This question is addressed in the application domain of sound design, where expert practitioners typically face the challenge of exploring high-dimensional, parametric sound spaces. We propose a user-centred design approach with expert sound designers to steer the design of such a system and better conceptualize exploration within this context.

5.1.1 Musical Task: Synthesis Exploration

Sound design is an exemplary application domain for studying exploration—taking iterative actions and multiple steps to move from an ill-formed idea to a concrete realization [Garcia et al., 2012]. Sonic exploration tasks can take myriad of forms: for example, composers explore various sketches of their musical ideas to write a final score; musicians explore different playing modes to shape an instrument's tone; sound designers explore several digital audio parameters to create unheard-of sounds [Monache et al., 2010, Delle Monache et al., 2018].

Situating the Musical Task

Most of today's digital commercial tools for sound synthesis, named Virtual Studio Technology (VST, see Fig. 5.2), still rely on complex interfaces using tens of technical parameters as inputs. These parameters often relate to the underlying algorithms that support sound synthesis, preventing users from establishing a direct perceptual relationship with the sound output. To that one may add the exponential number of parameter combinations, called presets, that eventually correspond to given sound designs. It is arguable that these interfaces may not be the best to support human exploration: as the perceptual outcome of acting on a given parameter may rapidly become unpredictable, they may hinder user appropriation [Resnick et al., 2005, Shneiderman, 2007].

As such, synthesis exploration is a general task that concern expert musicians—which explore parameters of computer devices to make sound—as well as non-

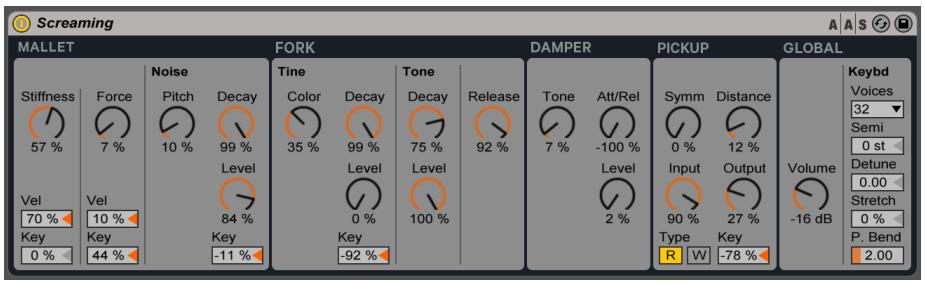


Figure 5.2: A typical VST interface for sound design, containing many technical parameters.

musicians—which explore parameters of any kind of interfaces when appropriating interaction. We are interested in investigating expert musicians' strategies to explore parametric interfaces. Our wish is that studying expert users will help us create new interaction techniques that may benefit non-expert users.

Human-Centred Approaches

Creativity support tools have long focused on exploration as a central task to human creative work [Shneiderman, 2007]. Design guidelines for supporting exploration were developed, which include aiming at simple interfaces for appropriating the tool and getting into sophisticated interaction more easily [Dix, 2007]. Flexible interaction modalities that can adapt to users' very own styles of thinking and creating may also be required [Resnick et al., 2005]. In particular, parameter space exploration remains a current challenge for HCI research [Cartwright et al., 2014]. Recently, creativity-oriented HCI researchers underlined the need to move toward interdisciplinary research collaborations [Frich et al., 2018].

Machine learning was in this sense examined for its implications in design [Koch, 2017] and identified as an opportunity for user experience [Dove et al., 2017, Yang et al., 2018a, Yang et al., 2018b]. Yet, a large body of work in the machine learning research community has so far focused on constructing autonomous algorithms learning creative behaviour from large amounts of impersonal data—falling under the name of computational creativity [Wiggins, 2006]. While this have allowed the building of powerful tools and models for creation, one may be concerned in the question of how to include human users in the design of such models to support human-computer co-creation [Kantosalo et al., 2014].

Davis et al. proposed a model of creativity that explicitly considers the computer as an enactive entity [Davis et al., 2014]. They notably stressed the potential of combining creativity support tools with computational creativity to enrich a collaborative process between the user and the computer [Davis et al., 2014]. The Drawing Apprentice, a co-creative agent that improvizes in real-time with users as they draw, illustrates their approach [Davis et al., 2016]. While their user study confirms the conceptual potential of building such artistic computer colleagues, its technical implementation remains specific to the use case at stake—*e.g.*, drawing. We propose

to jointly design a conceptual and technical framework that could be could easily be transferrable to other application domains—potentially realizing general mixed-initiative co-creativity [Horvitz, 1999, Yannakakis et al., 2014].

Machine-Centred Approaches

Interactive machine learning [Fails and Olsen Jr, 2003] allows human users to build customized models by providing their own data examples—typically a few of them. Not only users can customize training examples, but they are also allowed to directly manipulate algorithm parameters [Kapoor et al., 2010, Wong et al., 2011], as well as to receive information on the model's internal state [Amershi et al., 2015, Patel et al., 2011]. Applications in HCI cover a wide range of tasks, such as handwriting analysis [Shilman et al., 2006], recommender systems [Amershi et al., 2012], or prioritising notifications [Amershi et al., 2011]. Interactive machine learning mainly builds on supervised learning, which defines a computational framework for the learning of complex input-output models based on example input-output pairs. The "human-in-the-loop" approach to supervised learning critically differs from the computational creativity approach, which typically relies on huge, impersonal databases to learn models [Gillies et al., 2016].

Interactive machine learning is one such example of a generic framework for human-computer co-creation [Amershi et al., 2014]. The technical framework was successfully applied across several creative domains, such as movement interaction design [Zamborlin et al., 2014, Francoise and Bevilacqua, 2018, Gillies, 2019], web page design [Kumar et al., 2011] or video games [Kleinsmith and Gillies, 2013]. Specifically, research studying users building customized gestural controllers for music brought insight on the creative benefits of interacting with machine learning [Fiebrink et al., 2011]. Not only were users able to accomplish their design goal—*e.g.*, demonstrating a given gesture input for controlling a given sound parameter output—, but they also managed to explore and rapidly prototype alternative designs by structuring and changing training examples [Fiebrink et al., 2010]. These patterns were reproduced by novice users who gained accessibility using examples rather than raw parameters as input [Katan et al., 2015]. The algorithms' sometimes surprising and unexpected outcomes favoured creative thinking and sense of partnership in human users [Fiebrink and Caramiaux, 2016].

Typical workflows in interactive machine learning tend to iterate on designing training examples that are built from a priori representative features of the input space to support exploration. Yet, in some creative tasks where a problem definition may be found only by arriving at a solution [Dorst and Cross, 2001, Rittel, 1972], it might be unfeasible for users to define, a priori, such representative features of the final design [Katan et al., 2015]. Other approaches proposed methods to release such contraints, for example by exploring alternative machine learning designs by only defining the limits of some parameter space [Scurto et al., 2016]. We propose to further investigate machine learning frameworks able to iteratively learn from other user input modalities, and explicitly considering mixed-initiative workflows, where systems autonomously adapt to users [Deterding et al., 2017]. As reviewed in the next section, using interactive reinforcement learning offers such perspectives.

5.1.2 Proposed Technique: Deep Reinforcement Learning

Interactive reinforcement learning defines a computational framework for the interaction between a learning agent, a human user, and an environment [Amershi et al., 2014]. Specifically, users can communicate positive or negative *feedback* to the agent, in the form of a numerical reward signal, to teach it which action to take when in a certain environment state. The agent is thus able to adapt its behaviour to users, while remaining capable of behaving autonomously in its environment. Interactive reinforcement learning has been recently applied in HCI [Ruotsalo et al., 2014], with promising applications in exploratory search [Glowacka et al., 2013, Athukorala et al., 2016b] and adaptive environments [Frenoy et al., 2016, Rajaonarivo et al., 2017]. Integrating user feedback in reinforcement learning algorithms is computationally feasible [Stumpf et al., 2007], helps agents learn better [Knox and Stone, 2009], can make data-driven design more accessible [Lomas et al., 2016], and holds potential for rich human-computer collaboration [Stumpf et al., 2009]. Applications in Human-Robot Interaction informed on how humans may give feedback to learning agents [Thomaz and Breazeal, 2008], and showed potential for enabling human-robot co-creativity [Fitzgerald et al., 2017].

Deep Reinforcement Learning

Recently, reinforcement learning has witnessed a rise in popularity thanks to advances in deep neural networks [Mnih et al., 2015]. Powerful models including user feedback have been developed for high-dimensional parameter spaces [Christiano et al., 2017, Warnell et al., 2017]. Design researchers have identified reinforcement learning as a promising prospective technique to improve human-machine "joint cognitive and creative capacity" [Koch and Oulasvirta, 2018]. We believe that interactive reinforcement learning—especially deep reinforcement learning—holds great potential for supporting creative tasks—especially exploration of high-dimensional VST parameter spaces. First, its computational framework, constituted by environment states, agent actions, and user feedback, remains fully generic [Sutton and Barto, 2011], and thus potentially allow the design of generic interaction modalities transferrable to different application domains. Second, the autonomous behaviour intrinsic to reinforcement learning algorithms may be exploited to build a novel creative mixed-initiative paradigm, where the user and the agent would cooperate by taking actions that are "neither fully aligned nor fully in conflict" [Crandall et al., 2018]. Finally, we consider that user feedback could be a relevant input modality in the case of exploration, notably for expressing on-the-fly, arbitrary preferences toward imminent modifications, as opposed to representative examples. As previously stated, this requires investigating a somewhat unconventional use of reinforcement learning: if previous works employed user feedback to teach agents a "correct" behavior in relation to a task's goal, it is less obvious whether such a correct behavior may be well-defined—or even exists—for human users performing exploration.

Method

We adopted a user-centered approach to lead joint conceptual and technical work on interactive reinforcement learning for synthesis exploration. Two design iterations—a pilot study and an evaluation workshop—were conducted over the course of our research. Two prototypes were designed and developed—one initial reinforcement learning prototype, and the *Co-Explorer*, our final deep reinforcement learning prototype. The process thus includes sequentially:

- Prototype 1: Implementing a reinforcement learning algorithm that learns to explore VST parameter spaces from binary human feedback
- Pilot study: Observing and interviewing participants exploring sound spaces, first using standard parametric interfaces, then using our initial reinforcement learning prototype
- Prototype 2: Designing deep reinforcement learning in response to design ideas suggested by our pilot study
- Evaluation workshop: Observing and discussing with participants using and appropriating the *Co-Explorer*, our final prototype, in two creative tasks related to exploration

We worked with a total of 14 users (5 women, 9 men; all French) through the series of activities. From the 14 total, there were 2 who took part in all of the activities listed below, to testify of our prototype's improvements. Our users covered different areas of expertise in sound design and ranged from sound designers, composers, musicians, and artists to music researchers and teachers. Thus, they were not all constrained to one working methodology, one sonic practice or one application domain. Our motivation was to sample diverse approaches to exploration that sound design may provoke, in order to design a flexible reinforcement learning algorithm that may suit a variety of users' working styles [Resnick et al., 2005].

5.2 Prototyping I

By formalizing synthesis exploration as an interactive reinforcement learning problem, we seek to tackle both issues at once. First, human navigation in high-dimensional parameter spaces may be facilitated by the reinforcement learning computational framework, made of sequences of states, actions, and rewards. Second, human creativity may be stimulated by the autonomous behaviour of reinforcement learning algorithms, suggesting other directions or design solutions to users along exploration.

5.2.1 Algorithm: *Sarsa*

We used the same model prototype than that in Chapter 4, this time calling it "RL agent" for concision purposes. For the readers' convenience, we describe again its principles below.

The prototype lets users navigate through different sounds by only communicating positive or negative feedback to a reinforcement learning agent. The agent learns from feedback how to act on the underlying synthesis parameters in lieu of users (see Fig. 5.3). Formally, the environment is constituted by the VST parameters, and the agent iteratively acts on them. Computationally, we considered the state space $\mathcal{S} = \{S\}$ constituted by all possible parameter configurations $S = (s_1, ..., s_n)$, with n being the number of parameters, and $s_i \in [s_{min}, s_{max}]$ being the value of the i^{th} parameter living in some bounded numerical range (for example, s_i can control the level of noise normalized between 0 and 1). We defined the corresponding action space $\mathcal{A}(S) = \{A\}$ as moving up or down one of the n parameters by one step a_i, except when the selected parameter equals one boundary value:

$$A(S) = \begin{cases} \pm a_i & \text{for } s_i \in]s_{min}, s_{max}[\\ +a_i & \text{for } s_i = s_{min} \\ -a_i & \text{for } s_i = s_{max} \end{cases} \tag{5.1}$$

We implemented *Sarsa*, which is a standard algorithm to learn how to act in many different environment state, *i.e.,* for each given parameter configuration [Sutton and Barto, 2011]. It differs from multi-armed bandits, which learns how to act in one unique environment state [Lomas et al., 2016]. Importantly, as evoked in Section 5.1, *Sarsa* was designed to learn an optimal behaviour in relation to the goal of a task. Our purpose in this study was to scope the pros and cons of such a standard reinforcement learning algorithm for human exploration tasks, judging how it may influence user experience, and framing how it may be engineered with regard to this.

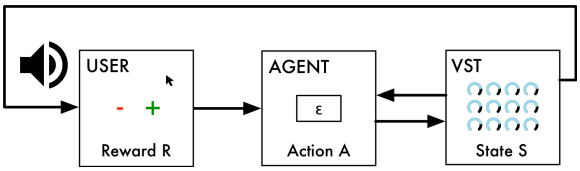


Figure 5.3: Graphical outline of our RL agent prototype. Users can only provide feedback to the agent, which acts on hidden VST parameters.

We used a VST-based 12-parameter space ($n = 12$) as the environment of our prototype. Because *Sarsa* is defined on discrete state spaces, each parameter range was discretized in three normalized levels ($s_i \in \{0, 0.5, 1\}, a_i = 0.5; 0 \leq i \leq n$). Although this would have been a design flaw in a perceptual experiment on typical VSTs, this allowed for obvious perceptual changes, which was required to investigate feedback-based interaction with a large variety of sounds.

An ε-greedy method defines the autonomous exploration behaviour policy of the agent—how it may act by exploiting its accumulated feedback while still exploring new unvisited states [Sutton and Barto, 2011]. It consists in having the agent take an optimal action with probability ε, and reciprocally, take a random action with probability $1 - \varepsilon$. For example, $\varepsilon = 1$ would configure an always exploiting agent—*i.e.,* always taking the best actions based on accumulated feedback—, while $\varepsilon = 0$ would

configure an always exploring agent—*i.e.,* never taking into account the received feedback. Our purpose in this study was to examine whether different exploration-exploitation trade-offs could map to different user approaches to exploration.

5.2.2 Human Data: Feedback

Finally, we propose that the user would be responsible for generating feedback. We directly mapped user feedback to the environmental reward signal R associated with a given state-action pair (S, A). The resulting formalization—where an agent takes actions that modifies the environment's state and learn from feedback received from a user—defines a generic interactive reinforcement learning problem.

5.3 Observation I

We organized a one-day pilot study with four of our expert participants. The aims of this pilot study were to: Observe approaches to exploration in standard VST interfaces; Identify problems users experience; Introduce the reinforcement learning technology in the form of a prototype; Brainstorm ideas and possible breakdowns.

5.3.1 Setup: Pilot Study

The pilot study was divided in two parts: (1) parametric interface exploration, then (2) interactive reinforcement learning-based exploration. We conducted individual semi-structured interviews at the end of each part, having each participant do the study one by one. This structure was intended to bring each participant to become aware of their subjective experience of exploration [Petitmengin, 2006]. Our intention was to open up discussions and let participants suggest design ideas about interactive reinforcement learning, rather than testing different algorithmic conditions in a controlled, experimental setup. We spent an average of 2 hours with each of our four participants, who covered different expertise in sound design (composition, sound design, interaction design, research).

5.3.2 Results: VST Interfaces, Qualitative Analysis

Procedure

In the first part of the study, participants were asked to find and create a sound preset of their choice using three different parametric interfaces with different number of parameters (respectively 2, 6, and 12, see Fig. 5.4). No reinforcement learning agent was used. We linked each interface to a different sound synthesis engine (respectively using FM synthesis[1], and one commercial VST from which we selected 6, then 12, parameters). Sound was synthesized continuously; participants' actions were limited to move the knobs using the mouse to explore the design space offered by all possible combinations. Knobs' technical names were hidden to test the generic effect of parameter dimensionality in interface exploration, and avoid any biases due to user

[1] *Frequency Modulation synthesis* (a classic algorithmic method for sound synthesis [Chowning, 1973]).

knowledge of parameter function (which typically occur with labelled knobs). Interface order was randomized; we let participants spend as much time as they wanted on each interface to let them explore the spaces freely.

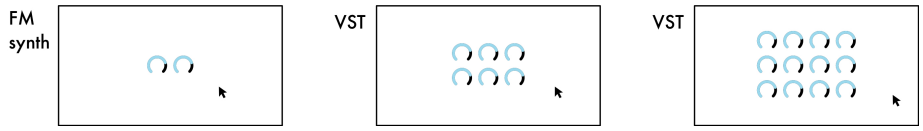


Figure 5.4: Graphical outline of the three parametric interfaces of our study.

Analysis

We were interested in observing potential user strategies in synthesis exploration. We thus logged parameter temporal evolution during the task. It consists in an n-dimensional vector, with n being the number of parameters (respectively 2, 6, then 12). Sample rate was set to 100 ms, which is a standard value for interaction with sound and musical interfaces [Jorda, 2005]. We used Max/MSP[2] and the MuBu[3] library to track user actions on parameters and record their evolutions. We used structured observation to study participants' interviews. This method was meant to provide a thorough qualitative analysis on user exploration strategies.

Observations

Qualitative analysis of parameter temporal evolution let us observe a continuum of approaches to parametric interface exploration. We call the first extremity of this continuum **analytical exploration**: this involves actioning each of the knobs one after the other over their full range. The second is called **spontaneous exploration**: this involves making random actions on the knobs. Figure 5.5 shows examples for each of these two approaches. One participant was consistently analytical over the three interfaces; one was consistently spontaneous over the three. The two others combined both approaches over the three interfaces.

Interview analysis let us map these approaches to different subgoals in exploration. The analytical approach concerns exploration of the interface at a parameter level: *"The strategy is to test [each knob] one by one to try to grasp what they do"*, one participant said. The goal of exploration is then related to building a mental map of the parameters to learn how to navigate in the design space. The spontaneous approach concerns exploration of the design space at a creative level: *"I moved the knobs more brutally and as a result of serendipity I came across into something different, that I preferred for other reasons..."*, another participant said. The goal of exploration is then related to discovering new parameter states leading to inspiring parts of the design space.

Discovery is critical to synthesis exploration. *"Once [the knobs] are isolated, you let yourself wander a bit more..."*, one participant analysed. Surprise is also important:

[2]https://cycling74.com/products/max/
[3]https://forum.ircam.fr/projects/detail/mubu/

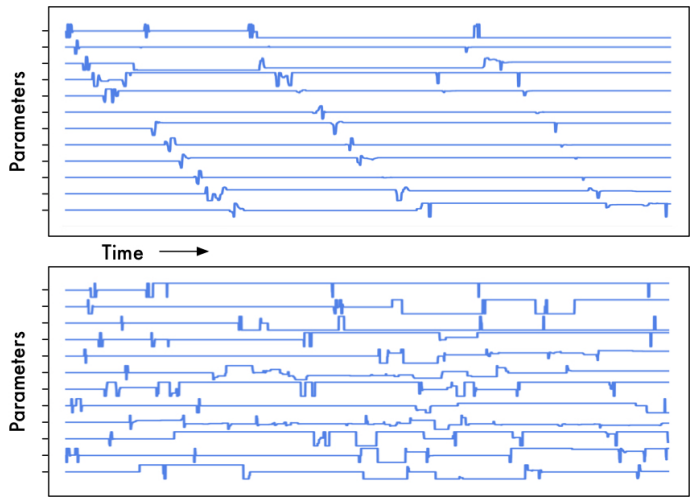


Figure 5.5: Two user exploration strategies with a 12-dimensional parametric interface: Analytical (top) vs. spontaneous (bottom).

"To explore is to be in a mental state in which you do not aim at something precise", one participant said. Interestingly, we observed that participants often used words related to perceptual aspects rather than technical parameters. *"I like when you can get a sound that is... um... Consistent, like, coherent. And at the same time, being able to twist in many different ways. This stimulates imagination, often"*, one participant said. Two participants mentioned that forgetting the parametric interface may be enjoyable in this sense: *"I appreciate an interface that does not indicate [...], that has you go back into sound, so that you are not here reading things, looking at symbols..."*, one participant said.

All participants reported being hindered in their exploration by the parameter inputs of the three interfaces. As expected, the more parameters the interface contained, the larger the design space was, and the harder it was to learn the interface. *"For me, the most important difficulty is to manage to effectively organise all things to be able to re-use them."*, one participant said. Time must be spent to first understand, then to memorize the role of parameters, taking into account that their role might change along the path of exploration. This hampers participants' motivation, often restraining themselves to a subspace of the whole design space offered by the tool: *"after a while I was fed up, so I threw out some parameters"*, one participant said about the 12-knob interface.

Participants discussed the limitations encountered in the study in light of their real-world practice with commercial interfaces. Two participants mentioned using automation functions to support synthesis exploration. Such functions include randomizing parameter values, automating parameter modification over time, or creating new control parameters that *"speak more to your sensibility, to your ears, than to what happens in the algorithm"*, to cite one of the participants. Two participants also use factory presets to start exploration: *"I think that in some interfaces they are pretty well*

conceived for giving you the basis of a design space. Then it's up to you to find what parameters to move", one participant said. Two participants said that the graphical user interfaces, including parameter names, knob disposition, and visual feedback on sound, may help them manage to lead exploration of large parameter spaces.

5.3.3 Results: RL Agent Prototype, Qualitative Analysis

Results in first part let us identify different user approaches to VST interface exploration, as well as different problems encountered in high-dimensional parameter spaces. In the second part, we were interested in having participants test the reinforcement learning technology in order to scope design ideas and possible breakthroughs in relation to exploration.

Procedure

Our participants were asked to find and create a sound preset of their choice by communicating feedback to three different agents with different exploration behaviours (respectively $\varepsilon = 0$; $\varepsilon = 1$; and $\varepsilon = 0.5$). Sound was synthesized continuously, in a sequential workflow driven by the agents' algorithmic functioning. At step t, participants could listen to a synthesized sound, and give positive or negative feedback by clicking on a two-button interface (Fig. 5.6). This would have the agent take an action on hidden VST parameters, modify the environment's state, and synthesize a new sound at step $t + 1$. Participants were only told to give positive feedback when the agent gets closer to a sound that they enjoy, and negative feedback when it moves away from it. They were not explained the agent's internal functioning, nor the differences between the three agents. The starting state for $t = 0$ was randomly selected. Agent order was randomized; we asked participants to spend between 5 and 10 minutes with each.



Figure 5.6: One of our four participants using a two-button interface to communicate binary feedback to the RL agent prototype in the pilot study.

Analysis

We logged all participant actions in the graphical user interface. It consisted in timed onsets for positive feedback on the one hand, and negative feedback on the other hand. We also logged parameter temporal evolution to observe how the RL agent would act on parameters following user feedback. We used structured observation to study participants' interviews and discussions led at the end of the pilot study.

Reactions

All participants reported forgetting synthesis parameters to focus on the generated sound. The simplicity and straightforwardness of the new interface benefited their exploration. *"There's always this sensation that finally you are more focused on listening to the sound itself rather than trying to understand the technology that you have under your hands, which is really great, yeah, this is really great"*, one participant said.

The computational framework defined by reinforcement learning was well understood by all participants. *"There's somewhat a good exploration design [sic], because it does a bit what you do [with the parametric interface], you move a thing, you move another thing..."*, one participant said. All participants enjoyed following agents' exploration behaviours, mentioning a playful aspect that may be useful for serendipity. Three participants in turn adapted their exploration to that of the agent: *"you convince yourself that the machine helps you, maybe you convince yourself that it is better... and after you go on exploring in relation to this"*, one participant said. Interestingly, one participant that was skeptical about partnering with a computer changed his mind interacting with the RL agent: *"We are all different, so are they"*, he commented, not without a touch of humor.

Uses of Feedback

Descriptive statistics informed on how participants used the feedback channel. Three participants gave feedback every 2.6 seconds on average ($\sigma = 0.4$), globally balancing positive with negative (average of 44.8% positive, $\sigma = 0.02$). The fourth participant gave feedback every 0.9 seconds on average ($\sigma = 0.07$) which was mostly negative (average of 17.2% positive, $\sigma = 0.02$). All participants reappropriated the feedback channel, quickly transgressing the task's instructions toward the two-button interface to fulfill their purposes. One participant used feedback to explore agents' possible behaviors: *"Sometimes you click on the other button, like, to see if it will change something, [...] without any justification at all"*, he commented. Another used the '-' button to tell the agent to *"change sound"*. Two participants also noticed the difference between feedback on sound itself, and feedback on the agent's behavior: *"there's the 'I don't like' compared to the sound generated before, and the 'I don't like it at all', you see"*, one of them said.

Breakdowns

Rapidly, though, participants got frustrated interacting with the RL agent. All participants judged that agents did not always reacted properly to their feedback, and were

leading exploration at the expense of them: *"sometimes you tell 'I don't like', 'I don't like', 'I don't like', but it keeps straight into it! (laughs)"*, one participant said. Contrary to what we expected, participants did not expressed a strong preference for any of the three tested agents. Only one participant noticed the randomness of the exploring agent, while the three other participants could not distinguish the three agents. This may be caused by the fact that the *Sarsa* algorithm was not designed for the interactive task of human exploration. Reciprocally, this may be induced by experiential factors due to the restricted interaction of our RL agent prototype, *e.g.*, preventing users to undo their last actions. Finally, two participants also complained about the lack of precision of the agent toward the generated sounds. This was induced by the *Sarsa* algorithm, which required to discretize the VST parameter space.

Design Implications

Participants jointly expressed the wish to lead agent exploration. They suggested different improvements toward our RL agent prototype:

- Express richer feedback to the agent (*e.g.*, differentiating *"I like"* from *"I really like"*)
- Control agent path more directly (*e.g.*, commanding the agent to go back to a previous state, or to some new unvisited state in the parameter space)
- Improve agent algorithm (*e.g.*, acting more precisely on parameters, reacting more accurately to feedback)
- Integrate agent in standard workspace (*e.g.*, directly manipulating knobs at times in lieu of the agent)

Interestingly, one participant suggested moving from current sequential workflow (where the agent waits for user feedback to take an action on the environment's state) to an autonomous exploration workflow (where the agent would continuously take actions on the environment's state, based on both accumulated and instantaneous user feedback). Three participants envision that such an improved RL agent could be useful in their practice, potentially allowing for more creative partnerships between users and agents.

5.4 Prototyping II

Our pilot study led us to the design of a final prototype, called *Co-Explorer*. We decided to first design new generic interaction modalities with RL agents, based on users' reactions with both parametric interfaces and our initial prototype.

5.4.1 Human Data: Feedback + Controls

Our initial prototype only employed user feedback as its unique interaction modality. This limited our participants, who suggested a variety of new agent controls to

support exploration. We translated these suggestions into new interaction modalities that we conceptualized under three generic categories: (1) user feedback, (2) state commands, and (2) direct manipulations (as shown in Fig. 5.7).

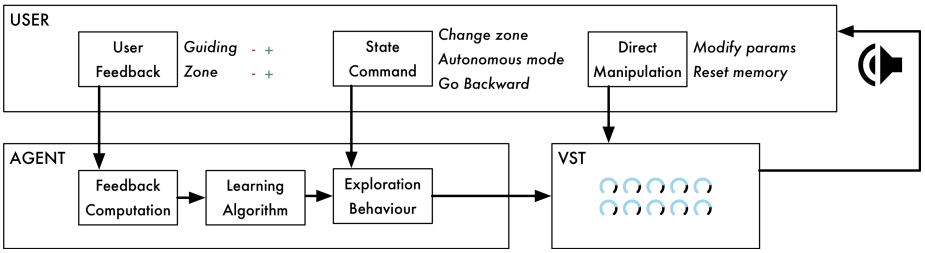


Figure 5.7: Interactive workflow of the *Co-Explorer*.

User Feedback

Our design intention is to support deeper user customization of the VST parameter space, while also allowing richer user contribution to agent learning. We thus propose to enhance user feedback as defined in our initial prototype, distinguishing between *guiding* and *zone feedback*. Guiding feedback corresponds to users giving binary guidance toward the agent's instantaneous trajectory in the parameter space. Users can give either positive—*i.e.,* "keep going in that direction"—or negative guidance feedback—*i.e.,* "avoid going in that direction". Zone feedback corresponds to users putting binary preference labels on given zones in the parameter space. It can either be positive—*i.e.,* "this zone interests me"—or negative—*i.e.,* "this zone does not interest me". Zone feedback would be used for making assertive customization choices in the design space, while guiding feedback would be used for communicating on-the-fly advice to the learning agent.

State Commands

Additionally, our design intention is to support an active user understanding of agent actions in the VST parameter space. We propose to define an additional type of interaction modality—we call them "state commands". State commands enable direct control of agent exploration in the parameter space, without contributing to its learning. We first allow users to command the agent to *go backward* to some previously-visited state. We also enable users to command the agent to *change zone* in the parameter space, which corresponds to the agent making an abrupt jump to an unexplored parameter configuration. Last but not least, we propose to let users start/stop an *autonomous exploration mode*. Starting autonomous exploration corresponds to letting the agent continuously act on parameters, possibly giving feedback throughout its course to influence its behaviour. Stopping autonomous exploration corresponds to going back to the sequential workflow implemented in our initial prototype, where the agent waits for user feedback before taking a new action on parameters.

Direct Manipulation

Lastly, our design intention is to augment, rather than replace, VST interfaces with interactive reinforcement learning, leveraging users expertise with these interfaces and providing them with additional modalities that they could solicit when they may need it. We thus propose to add "direct manipulations" to support direct parameter modification through a standard parametric interface. It lets users explore the space on their own by only manipulating parameters without using the agent at all. It can also be used to take the agent to a given point in the VST parameter space—*i.e.,* "start exploration from this state"—, or to define by hand certain zones of interest using a zone feedback—*i.e.,* "this example preset interests me". Inversely, the parametric interface also allows to visualize agent exploration in real-time by observing how it acts on parameters.

A last, global interaction modality consists in *resetting agent memory*. This enables users to start exploration from scratch by having the agent forget accumulated feedback. Other modalities were considered, such as modifying the agent's speed and precision. Preliminary tests pushed us to decide not to integrate them in the *Co-Explorer*.

5.4.2 Algorithm: *Deep TAMER*

Based on our observations in the pilot study, we developed our reinforcement learning agent at two intertwined technical levels: (1) feedback formalization and (2) learning algorithm. This work was done in the context of Bavo Van Kerrebroeck's Master's Thesis, supervised by Frédéric Bevilacqua and myself [Van Kerrebroeck, 2018].

Feedback Formalization

One challenge consisted in addressing the non-stationarity of user feedback data along their exploration. We implemented *Deep TAMER*, a reinforcement learning algorithm suited for human interaction [Warnell et al., 2017]. *Deep TAMER* leverages a feedback formalization that distinguishes between the environmental reward signal—*i.e.,* named R in the *Sarsa* algorithm of our initial prototype—and the human reinforcement signal—*e.g.,* feedback provided by a human user. This technique, implemented in the *TAMER* algorithm [Knox and Stone, 2009], was shown to reduce sample complexity over standard reinforcement learning agents, while also allowing human users to teach agents a variety of behaviours. We detail the differences between standard RL algorithms and (deep) TAMER in Appendix A.

Learning Algorithm

Another challenge was to tackle learning in high-dimensional parametric spaces that are typical of our use case. *Deep TAMER* employs function approximation [Sutton and Barto, 2011] to generalize user feedback given on a subset of state-action pairs to unvisited state-action pairs. Specifically, a deep neural network is used to learn the best actions to take in a given environment state, by predicting the amount of user feedback it will receive [Mnih et al., 2015, Warnell et al., 2017]. The resulting algorithm can

learn in high-dimensional state spaces $\mathcal{S} = \{S\}$ and is robust to changes in discretization a_i of the space. For our application in sound design, we engineered the algorithm for $n = 10$ parameters. We normalized all parameters and set the agent's precision by discretizing the space in one hundred levels ($s_i \in [0, 1], a_i = 0.01; 0 \leq i \leq n$).

A last challenge was to learn quickly from the small amounts of data provided by users during interaction. *Deep TAMER* uses a replay memory, which consists in storing the received human feedback in a buffer $\mathcal{D}$, and sampling repeatedly from this buffer with replacement [Warnell et al., 2017]. This was shown to improve the learning of the deep neural network in high-dimensional parameter spaces in the relatively short amount of time devoted to human interaction. We set the parameters of the the deep neural network by performing a parameter sweep and leading sanity checks with the algorithm.

5.5 Engineering II

We then engineered our model prototype by (1) improving exploration behaviour, (2) processing user feedback and controls, and (3) implementing a new specific interface for sound design. This work was done in the context of Bavo Van Kerrebroeck's Master's Thesis, supervised by Frédéric Bevilacqua and myself [Van Kerrebroeck, 2018].

5.5.1 Algorithm: Improving Parameter Exploration

We developed a novel exploration method for autonomous exploration behaviour. It builds on an intrinsic motivation method, which pushes the agent to "explore what surprises it" [Bellemare et al., 2016]. Specifically, it has the agent direct its exploratory actions toward uncharted parts of the space, rather than simply making random moves—as in the ε-greedy approach implemented in our initial prototype. It does so by building a density model of the parameter space based on all visited states. We used tile coding, a specific feature representation extensively used in the reinforcement learning literature to efficiently compute and update the density model in high-dimensional spaces [Sutton and Barto, 2011]. We parameterized ε with an exponential decay in such a way that its initial value would slowly decrease along user exploration. For our application in sound design, agent speed in continuous exploration mode was set to one action by tenths of a second.

5.5.2 Human Data: Processing Feedback and Controls

To fully realize our interaction design, we integrated the modalities defined in Section 5.4.1 within the reinforcement learning framework defined in Section 5.4.2.

User Feedback

We developed generic methods corresponding to user feedback modalities defined in Section 5.4.1 that we used in the feedback formalization of Section 5.4.2. For guiding feedback, we assigned user positive or negative feedback value over the p last state-action pairs taken by the agent (see Fig. 5.9, left), with a decreasing credit given

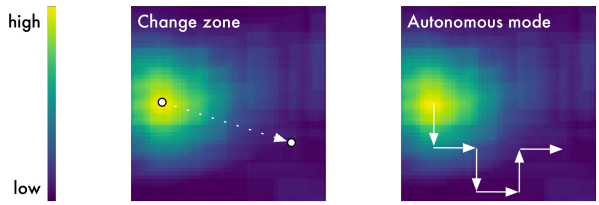


Figure 5.8: Graphical outline of machine exploration methods. The color scale depicts the density model all states. Left: Changing zone has the agent jump to the state with lowest density. Right: Autonomous exploration has the agent take successive actions leading to the state with lowest density.

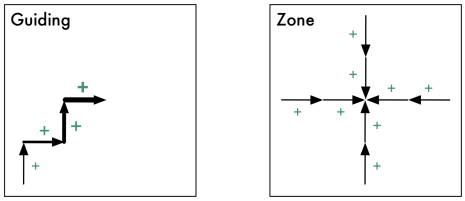


Figure 5.9: Graphical outline of human feedback computation methods. Here, positive feedback is given in some state situated at the center of the square. Left: Guiding feedback is distributed over the p lastly-visited state-action pairs. Right: Zone feedback impacts all state-action pairs potentially leading to the labelled state.

by a Gamma distribution [Knox and Stone, 2009]. For zone feedback, we computed all possible state-action pairs leading to the state being labelled and impacted them with positive or negative feedback received (see Fig. 5.9, right). This enables to build attractive and repulsive zones for the agent in the parameter space. Finally, we added a reward bonus to user feedback to enhance the agent's learning relatively to the novelty of a state. This reward bonus is computed using the density model described in Section 5.5.1.

State Commands

We developed generic methods corresponding to state commands defined in Section 5.4.1 using the exploration behaviour defined in Section 5.5.1. Changing zone has the agent randomly sampling the density distribution and jump to the state with lowest density (see Fig. 5.8, left). Autonomous exploration mode has the agent take exploratory actions that lead to the nearest state with lowest density with probability ε (see Fig. 5.8, right).

Direct Manipulation

We integrated direct manipulations as defined in Section 5.4.1 by leveraging the learning algorithm defined in Section 5.4.2. When parameters are modified by the user, the

reinforcement learning agent converts all parameters' numerical values as a state representation, taking advantage of the algorithm's robustness in changes of discretization. Reseting agent memory has the reinforcement learning algorithm erase all stored user feedback and trajectory, and load a new model.

5.5.3 Implementation: *Co-Explorer*

We implemented all interaction modalities described above in the *Co-Explorer* software (see Video for an overview).

Agent

We implemented the *Co-Explorer* as a Python library[4]. It allows to connect the deep reinforcement learning agent to any external input device and output software, using the OSC protocol for message communication [Wright, 2005]. This was done to enable future applications outside the sound design domain. Each of the features described in Section 5.4.2 are implemented as parameterized functions, which supports experimentation of interactive reinforcement learning with various parameter values as well as order of function calls. The current version relies on TensorFlow for deep neural network computations. The complete algorithm implementation and all learning parameters are shown in the Appendix.

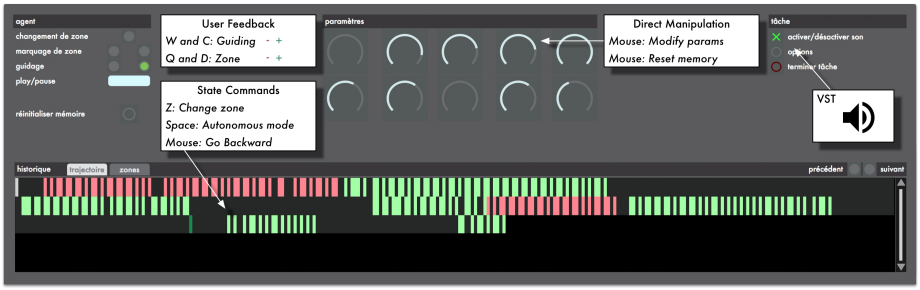


Figure 5.10: User interface of the *Co-Explorer*.

Interface

We implemented an interactive interface for our application in sound design (Fig. 5.10), which integrates all interaction modalities defined in Section 5.4.1. It builds on Max/MSP, a visual programming environment for real-time sound synthesis and processing. Standard parametric knobs enable users to directly manipulate parameters, as well as to see the agent act on it in real-time. An interactive history allows users to command the agent to go to a previously-visited state, be they affected by user feedback (red for negative, green for positive) or simply passed through (grey). Keyboard inputs support user feedback communication, as well as state commands that control agent exploration (changing zone, and start/stop continuous exploration mode). Lastly, a clickable button enables users to reset agent memory.

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

5.6 Observation II

We evaluated interaction with the *Co-Explorer* in a workshop with a total of 12 professional users (5 female, 7 male). The aims of the workshop were to: Evaluate each interaction modality at stake in the *Co-Explorer*; understand how expert users may appropriate the agent to support synthesis exploration.

5.6.1 Setup: Workshop

The workshop was divided in two tasks: (1) explore to discover, and (2) explore to create. This structure was intended to test the *Co-Explorer* in two different creative tasks (described in Section 5.6.2 and 5.6.3, respectively). Participants ranged from sound designers, composers, musicians, and artists to music researchers and teachers. They were introduced to the agent's interactive modalities and its internal functioning at the beginning of the workshop. In each part, they were asked to report their observations by filling a browser-based individual journal. Group discussion was carried on at the end of the workshop to let participants exchange views over synthesis exploration. The workshop lasted approximately three hours each.



Figure 5.11: Picture of sound designers using the *Co-Explorer* in our workshop.

5.6.2 Results: Discovery Task, Qualitative Analysis

Procedure

In the first part of the workshop, participants were presented with one VST (see Fig. 5.11). They were asked to use the *Co-Explorer* to explore and discover the sound space at stake. Specifically, we asked them to find and select five presets to constitute a rep-

resentative sample of the space. We defined the parameter space by selecting ten parameters from a commercial VST. Participants were encouraged to explore the space thoroughly. The task took place after a 10-minute familiarizing session: individual exploration lasted 25 minutes, followed by 5 minutes of sample selection, and 20 minutes of group discussion.

Analysis

All participant's actions were logged into a file. These contained timed onsets for user feedback—*i.e.,* binary guiding and zone feedback—, state commands—*i.e.,* backward commands in the history, changing zone commands, and autonomous exploration starting/stopping—, and direct manipulations—*i.e.,* parameter temporal evolutions. We also logged timed onsets for preset selection in relation to the task, but did not include the five presets themselves into our analysis. Our motivation was to focus on the process of exploration in cooperation with the *Co-Explorer*, rather than on the output of it. We used structured observation to extract information from individual journals and group discussion.

Results

We first looked at how users employed state commands. Specifically, the autonomous exploration mode, which consisted in letting the agent act cotinuously on parameters on its own, was an important new feature compared to our sequentiam initial RL agent prototype. Participants spent more than half of the task using the *Co-Explorer* in this mode (total of 13 minutes on average, $\sigma = 4.7$). Ten participants used autonomous exploration over several short time slices (average of 50 seconds, $\sigma = 25$s), while the two remaining participants used it over one single long period (respectively 9 and 21 minutes). P5 commented about the experience: *"It created beautiful moments during which I really felt that I could anticipate what it was doing. That was when I really understood the collaborative side of artificial intelligence"*.

The changing zone command, which enabled to jump to an unexplored zone in the parameter space, was judged efficient by all participants to find diverse sounds within the design space. It was used between 14 and 90 times, either to start a new exploration (P1: *"Every time I used it, I found myself in a zone that was sufficiently diametrically opposed to feel that I could explore something relatively new"*), or to rapidly seize the design space in the context of the task (P12: *"I felt it was easy to manage to touch the edges of all opposite textures"*). Interestingly, P2 noticed that the intrisic motivation method used for agent exploration behaviour *"brought something more than a simple random function that is often very frustrating"*.

We then looked at how users employed feedback. Guiding feedback was effectively used in conjunction with autonomous exploration by all participants, balancing positive with negative (55% positive on average, $\sigma = 17\%$). Participants gave various amounts of guiding feedback (between 54 and 1489 times). These strategies were reflected by different reactions toward the *Co-Explorer*. For example, one participant was uncertain in controlling the agent through feedback: *"if the agent goes in the right direction, I feel like I should take time to see where it goes"*, he commented. On

the contrary, P1 was radical in his controlling the agent, stating that he is *"just looking for another direction"*, and that he uses feedback *"without any value judgement"*. This reflects the results described in Section 5.3.3 using our initial RL agent prototype.

Zone feedback, enabling customization of the space with binary labels, was mostly given as positive by participants (72%, $\sigma = 18\%$). Two participants found the concept of negative zones to be counter-intuitive. *"I was a bit afraid that if I label a zone as negative, I could not explore a certain part of the space"*, P8 coined. This goes in line with previous results on applying interactive reinforcement learning in the field of robotics [Thomaz and Breazeal, 2008]. All participants agreed on the practicality of combining positive zone feedback with backward state commands in the history to complete the task. *"I labeled a whole bunch of presets that I found interesting [...] to after go back in the trajectory to compare how different the sounds were, and after continue going in other zones. I found it very practical"*, P8 reported. Overall, zone feedback was less times used than guiding feedback (between 10 and 233 times).

Finally, direct manipulation was deemed efficient by participants in certain zones of the design space. *"When I manage to hear that there is too much of something, it is quicker to parametrize sound by hand than to wait for the agent to find it itself, or to learn to detect it"*, P4 analyzed. P10 used them after giving a backward state command, saying she *"found it great in cases where one is frustrated not to manage to guide the agent"*. P11 added that she directly manipulate parameters to *"adjust the little sounds that [she] selected"*. P1 suggested that watching parameters move as the agent manipulates them could help learn the interface: *"From a pedagogical point of view, [the agent] allows to access to the parameters' functioning and to the interaction between these parameters more easily [than without]"*. This supports the fact that machine learning visualizations may be primordial in human-centred applications to enable interpretability of models [Amershi et al., 2014].

Relevance to Task

Three participants complained that the *Co-Explorer* did not react sufficiently quickly to feedback in relation to the task: *"I would really like to feel the contribution of the agent, but I couldn't"*, P12 said. Also, P3 highlighted the difficulties to give evaluative feedback in the considered task: *"without a context, I find it hard"*, he analysed. Despite this, all participants wished to spend more time teaching the *Co-Explorer*, by carefully customizing the parameter space with user feedback. For example, five participants wanted to slow the speed of the agent during autonomous exploration to be able to give more precise guidance feedback. Also, three participants wanted to express sound-related feedback: *"There, I am going to guide you about the color of the spectrum. [...] There, I'm going to guide you about, I don't know, the harmonic richness of the sound, that kind of stuff..."*, P4 imagined.

5.6.3 Results: Creation Task, Qualitative Analysis

Procedure

In the second part of the workshop, participants were presented with four pictures (Fig. 5.12). For each of these four pictures, they were asked to explore and create two sounds that subjectively depict the atmosphere of the picture. In this part, we encouraged participants to appropriate interaction with the *Co-Explorer* and feel free to work as they see fit. We used a new sound design space for this second part, which we designed by selecting another ten parameters from a commercial VST. Individual exploration and sound selection lasted 30 minutes, followed by 20 minutes of group discussion and 10 minutes of closing discussion.

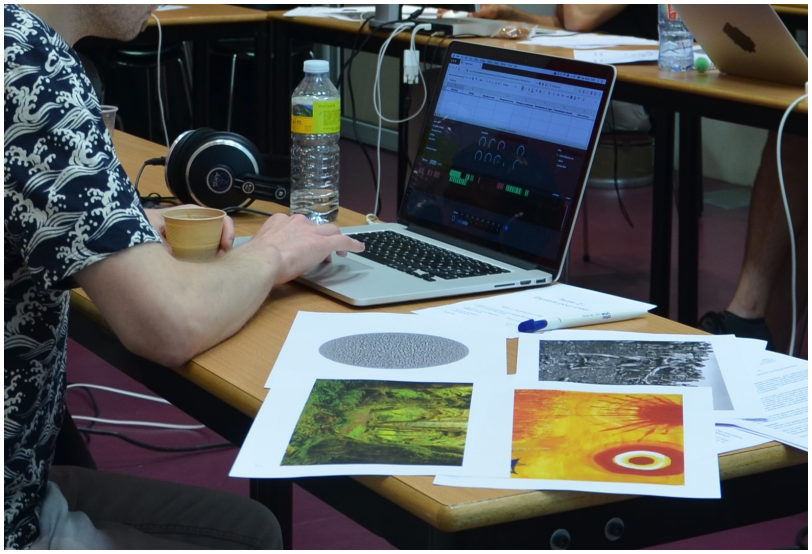


Figure 5.12: The four pictures framing the creation task of the workshop.

Analysis

All participant actions were logged into a file, along with timed parameter presets selected for the four pictures. Again, we focused our analysis on the process of exploration rather than on the output of it. Specifically, for this open-ended, creative task, we did not aim at analysing how each agent interaction modality individually relates to a specific user intention. Rather, we were interested in observing how users may appropriate the mixed-initiative workflow at stake in the *Co-Explorer*.

We used Principal Component Analysis (PCA [Jolliffe, 2011]), a dimensionality reduction method, to visualize how users switched parameter manipulation with agents. We first concatenated all participants' parameter evolution data as an n-dimensional vector to compute the two first principal components. We then projected each participant data onto these two components to support analysis of each user trajectory on a common basis. By doing this, relatively distant points would correspond to abrupt

changes made in parameters (*i.e.*, to moments when the user takes the lead on exploration). Continuous lines would correspond to step-by-step changes in parameters (*i.e.*, to moments when the *Co-Explorer* explores autonomously). PCA had a stronger effect in the second part of our workshop. We interpret this as a support to the two-part structure that we designed for the workshop, and thus did not include analysis of the first part. Finally, we used structured observation to extract information from individual journals and group discussion.

Exploration Strategies

All participants globally expressed more ease interacting with the *Co-Explorer* in this second task. *"I felt that the agent was more adapted to such a creative, subjective... also more abstract task, where you have to illustrate. It's less quantitative than the first task"*, P9 analysed. User feedback was also reported to be more intuitive when related to a creative goal: *"all parameters took their sense in a creative context. [...] I quickly found a way to work with it that was very efficient and enjoyable"*, P5 commented. Figure 5.13 illustrates the PCA for two different users interacting with the *Co-Explorer*.

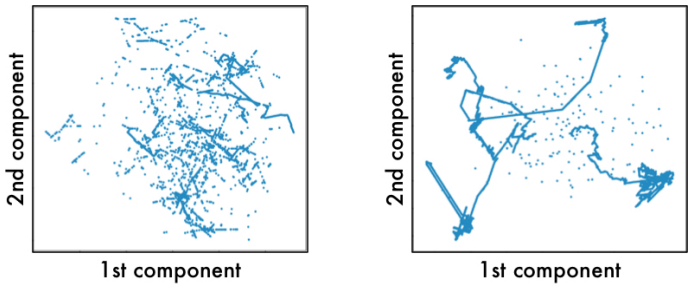


Figure 5.13: Two types of co-exploration partnerships shown in PCA visualizations of parameter evolution: User-as-leader (P9, left) and agent-as-leader (P7, right). Relatively distant points correspond to abrupt changes made in parameters (*i.e.*, to moments when the user takes the lead). Continuous lines correspond to step-by-step changes in parameters (*i.e.*, to moments when the *Co-Explorer* takes the lead).

Qualitative analysis of PCAs let us conceptualize a continuum of partnerships between our participants and the *Co-Explorer*. These could be placed anywhere between the two following endpoints:

- **User-as-leader**: This typically involves users first building a map of the design space (iteratively using changing zone and positive zone feedback), then generating variations of these presets (either through direct manipulation or short autonomous explorations).
- **Agent-as-leader**: This typically involves letting the *Co-Explorer* lead parameter manipulation (using autonomous exploration and guiding feedback), first setting some starting point in the design space (either using changing zone or direct manipulation).

Our interpretation is as follows. User-as-leader partnership may correspond to user profiles that approach creative work as a goal-oriented task, where efficacy and control are crucial (P10: *"I am accustomed... Where I work, if you prefer, we have to get as quick as possible to the thing that works the best, say, and I cannot spend so much time listening to the agent wandering around"*). Reciprocally, agent-as-leader partnership may correspond to user profiles that approach creative work as an open-ended task, where serendipity is essential for inspiration (P5: *"I did not try to look for the sound that would work the best. I rather let myself be pushed around, even a bit more than in my own practice"*). Some participants did not stabilize into one single partnership, but rather enjoyed the flexibility of the agent. *"It was quite fun to be able to let the agent explore, then stop, modulate a bit some parameters by hand, let it go and guide it again, changing zones too, then going back in the history... Globally, I have the impression of shaping, somewhat... I found it interesting"*, P11 coined.

Agent memory was handled with relevance to various creative processes toward the pictures. Seven participants disposed all four pictures in front of them (P7: *"to always have them in mind. Then, depending on the agent's exploration, I told myself 'hey, this sound might correspond to this picture'"*). Three participants focused on one picture at a time, *"without looking at the others"*. Four participants never reset the memory (P11: *"my question was, rather, in this given sonic landscape, how can I handle these four pictures, and reciprocally"*), and three participants reset agent memory for each of the different atmospheres shared by the pictures. Overall, participants benefited from partnering with the *Co-Explorer* in synthesis exploration: *"It's a mix of both. I easily managed to project a sound on the picture at first glance, then depending on what was proposed, it gave birth to many ideas"*, one participant said.

Toward Real-World Usages

All participants were able to describe additional features for the *Co-Explorer* to be usable in their real-world professional work environments—examples are, among others, connection to other sound spaces, memory transfer from one space to another, multiple agent memory management, or data exportation. They also anticipated creative uses for which the *Co-Explorer* were not initially designed. Half of the participants were enthusiastic about exploiting the temporal trajectories as actual artifacts of their creation (P6: *"What I would find super interesting is to be able to select the sequences corresponding to certain parameter evolution, or playing modes. [...] It would be super great to select and memorize this evolution, rather than just a small sonic fragment"*). Finally, two participants further imagined the *Co-Explorer* to be used as musical colleagues—either as improvisers with which one could *"play with both hands"* (P2), or as *"piece generators"* (P6) themselves.

5.7 Discussion

Our process of research, design, and development led to contributions at three different levels: (1) conceptual insight on human exploration; (2) technical insight on reinforcement learning; and (3) joint conceptual and technical design guidelines on

machine learning for creative applications.

5.7.1 Conceptual Insight

From Exploration to Co-Exploration

Our work with interactive reinforcement learning allowed for observing and characterizing user approaches to synthesis exploration, and supported it. While manipulating unlabelled parametric knobs of sound synthesizers, participants alternated between an *analytical* approach—attempting to understand the individual role of each parameter—and a *spontaneous* approach that could lead to combinations in the parameter space that might not be guessed with the analytical approach. While interacting with a reinforcement learning agent, participants tended to alternate the lead in new types of mixed-initiative workflows [Horvitz, 1999] that we propose to call *co-exploration* workflows. *User-as-leader* workflow was used for gaining control over each parameter of the design space. *Agent-as-leader* workflow allowed to relax users' control and provoke discoveries through the specific paths autonomously taken by the agent in the parameter space. Importantly, the benefit of interactive reinforcement learning for co-exploring sound spaces was dependent on the task. We found that this co-exploration workflow were more relevant to human exploration tasks that have a focus on creativity, such as in our workshop's second task, rather than discovery. Therefore, we believe that this workflow is well-suited in cases where exploration is somehow holistic (as in the creative task) rather than analytic (as in the discovery task where the goal is to understand the sound space to find new sounds).

Methodology

Our user-centered approach to exploration with interactive reinforcement learning allowed us to rapidly evaluate flexible interaction designs without focusing on usability. This process let us discover innovative machine learning uses that we may not have anticipated if we had started our study with an engineering phase. The simple, flexible, and adaptable designs tested in our first pilot study (parametric vs. RL) could in this sense be thought as technology probes [Hutchinson et al., 2003]. Working with professional users of different background and practices—from creative coders to artists less versed in technology—was crucial to include diverse user feedback in the design process. Our results support this, as many user styles were supported by the *Co-Explorer*. That said, user-driven design arguably conveys inherent biases of users. This is particularly true when promoting AI in interactive technology [Amershi et al., 2019, Caramiaux et al., 2019]. As a matter of fact, alongside a general enthusiasm, we did observe a certain ease among our professional users for expressing tough critiques, at times being skeptical on using AI, especially when the perception of the algorithm choice would contradict their spontaneous choice. Yet, the two professional users that took part to both our pilot study and workshop found the use of AI as welcome, testifying of its improvement along the development process.

Evaluation

Lastly, evaluation of reinforcement learning tools for creativity remains to be investigated more deeply. While our qualitative approach allowed us to harvest thoughtful user feedback on our prototypes' interaction modalities, it is still hard to account for direct links between agent computations and user creative goals. Using questionnaire methods, such as the Creativity Support Index [Cherry and Latulipe, 2014], may enable to measure different dimensions of human creativity in relation to different algorithm implementations. Also, focusing on a specific user category could also allow more precise evaluation in relationship to a situated set of creative practices and uses. Alternatively, one could aim at developing new reinforcement learning criteria that extends standard measures—such as convergence or learning time [Sutton and Barto, 2011]—to the qualitative case of human exploration. Research on interactive supervised learning has shown that criteria usually employed in the field of Machine Learning may not be adapted to users leading creative work [Fiebrink et al., 2011]. We believe that both HCI and ML approaches may be required and combined to produce sound scientific knowledge on creativity support evaluation.

5.7.2 Technical Insight

Computational Framework

Our two working prototypes confirmed that interactive reinforcement learning may stand as a generic technical framework for synthesis exploration. The computational framework that we proposed in Section 5.2.1, leveraging states, actions, and rewards, strongly characterized the mixed-initiative co-exploration workflows observed in Section 5.6.3—*e.g.,* making small steps and continuous trajectories in the parameter space. Other interactive behaviours could have been implemented—*e.g.,* allowing the agent to act on many parameters in only one action, or using different a_i values for different action sizes—to allow for more diverse mixed-initiative behaviours. Alternatively, we envision that domain-specific representations may be a promising approach for extending co-exploration. In the case of sound design, one could engineer high-level state features based on audio descriptors [Schwarz and Schnell, 2009] instead of using raw parameters. This could allow RL agents to learn state-action representations that would be independent from the parameter space explored—potentially allowing memory transfer from one parameter state space to another. This could also enable agent adaptation of action speed and precision based on perceptual features of the parameter space—potentially avoiding abrupt jumps in sound spaces.

Learning Algorithm

Reinforcement learning algorithmic functioning, enabling agents to learn actions over states, was of interest for our users, who were enthusiastic in teaching an artificial agent by feedback. Our deep reinforcement learning agent is a novel contribution to HCI research compared to multi-armed bandits (which explore actions over one unique state [Lomas et al., 2016]), contextual bandits (which explore in lower-dimensional state spaces [Koch et al., 2019]), and bayesian optimization (which ex-

plores at implicit scales [Shahriari et al., 2016]). We purposely implemented heterogeneous ways of teaching with feedback based on our observations of users' approaches to synthesis exploration, which extends previous implementations such as those in the Drawing Apprentice [Davis et al., 2016]. Yet, rich computational models of user feedback for exploration tasks remain a challenge. Our observations indeed suggested that exploring users may not generate a goal-oriented feedback signal, but may rather have several sub-optimal goals. They may also make feedback mistakes, act socially toward agents, or even try to trigger surprising agent behaviours over time. *Deep TAMER* was adapted to the interactive of user feedback (as opposed to *Sarsa*); yet, it still made the assumption that users will generate a stationary and always correct feedback signal [Warnell et al., 2017]. Previous works investigating how users give feedback to machine learning [Stumpf et al., 2009] may need to be extended to include such creative use cases.

Exploration Behaviours

The exploration behaviours of reinforcement learning agents were shown promising for fostering creativity in our users. Both ε-greedy and intrisic method were adapted to the interactive case of a user leading exploration. One of our users felt that intrisic motivation had agents behave better than random. Yet, users' perception of agent exploration behaviours remains to be investigated more deeply. In a complementary work [Scurto et al., 2018a], we confirmed that users perceived the difference between a random parameter exploration and a RL agent exploration. Yet, they might not perceive the difference between various implementations of agent exploration; what they perceive may be more related to the agent's global effect in exploring the parameter space. Future work may study co-exploration partnerships over longer periods of time to inquire co-adaptation between users and agents [Mackay, 1990]. On the one hand, users could be expected to learn to provide better feedback to RL agents to fulfill their creative goals—as it was shown in interactive approaches to supervised learning [Fiebrink et al., 2011]. On the other hand, agents could be expected to act more in line with users by exploiting larger amounts of accumulated feedback data—as it is typical with interactive reinforcement learning agents [Sutton and Barto, 2011]. A more pragmatic option would be to give users full control over agent epsilon values—*e.g.*, using an interactive slider [Koch et al., 2019]—to improve partnership in this sense.

5.7.3 Guidelines for Designing Machine Learning for Human Creativity

Based on our work with reinforcement learning, we identified a set of design challenges for leading joint conceptual and technical development of other machine learning frameworks for creative HCI applications. We purposely put back quotes from our participants in this section to inspire readers with insights on AI from users outside our design team.

Engage Users with Machine Learning

The *Co-Explorer* enabled users to fully engage with reinforcement learning computational framework. Users could explore as many states, provide as much feedback, and generate as many agent actions as they wanted to. They also had access to agent memory, be it by navigating in the interactive history, or by reseting the learned behaviour. In this sense, they had full control over the algorithmic learning process of the agent. This is well articulated by a participant, whose quote can be reported here: *"I did not feel as being an adversary to, or manipulated, by the system. A situation that can happen with certain audio software that currently use machine learning, where it is clear that one tries to put you on a given path, which I find frustrating—but this was not the case here"*.

These observations suggest that user engagement at different levels of machine learning processes may be essential to create partnering flows [Pachet et al., 2013]. That is, users should be provided with interactive controls and simple information on learning to actively direct co-creation. This goes in line with previous works studying user interaction with supervised learning in creative tasks [Amershi et al., 2014], which showed how users can build better partnerships by spending time engaging with algorithms [Fiebrink et al., 2011]. Careful interaction design must be considered to balance full automation with full user control and aim at creating flow states among people [Csikszentmihalyi, 1997]. Aiming at such user engagement may also constitute a design opportunity to demystify AI systems, notably by having users learn from experience how algorithms work with data [Fiebrink, 2019].

Foster Diverse Creative Processes

Our work showed that the *Co-Explorer* supported a wide diversity of creative user processes. Users could get involved in open-ended, agent-led exploration, or decide to focus on precise, user-led parameter modification. Importantly, none of these partnerships were clearly conceptualized at the beginning of our development process. Our main focus was to build a reinforcement learning agent able to learn from user feedback and to be easily controllable by users. In this sense, the *Co-Explorer* was jointly designed and engineered to ensure a dynamic human process rather than a static media outcome. As a matter of fact, we report one participant's own reflection, which we believe illustrate our point: *"What am I actually sampling [from the parameter space]? Is is some kind of climate that is going to direct my creation afterwards? [...] Or am I already creating?"*.

This suggests that supporting the process of user appropriation may be crucial for building creative AI partnerships. Many creative tools based on machine learning often focus on engineering one model to ensure high performance for a given task. While these tools may be useful for creative tasks that have a focus on high productivity, it is arguable whether they may be suited to creative work that has a focus on exploration as a way to build expression. For the latter case, creative AI development should not focus on one given user task, but should rather focus on providing users with a dynamic space for expression allowing many styles of creation [Resnick et al., 2005]. The massive training datasets, which are usually employed in the Ma-

chine Learning community to build computational creativity tools, may also convey representational and historical biases among end users [Suresh and Guttag, 2019]. Interactive approaches to machine learning directly address this issue by allowing users to intervene in real-time in the learning process [Fiebrink and Caramiaux, 2016].

Steer Users Outside Comfort Zones

The *Co-Explorer* actively exposed the exploration behaviour of reinforcement learning to users. This goes in opposition with standard uses of these algorithms [Brockman et al., 2016], and may provoke moments where agents behaviours may not align with users creative drive [Crandall et al., 2018]. Yet, it managed to build *"playful"* and *"funny"* partnerships that led some users to reconsider their approach to creativity, as one participant confessed: *"At times, the agent forced me to try and hear sounds that I liked less—but at least, this allowed me to visit unusual spaces and imagine new possibilities. This, as a process that I barely perform in my own creative practice, eventually appeared as appealing to me".*

This suggests that AI may be used beyond customisation aspects to steer users outside their comfort zones in a positive way. That is, designers should exploit non-optimal algorithmic behaviours in machine learning methods to surprise, obstruct, or even challenge users inside their creative process. Data-driven user adaptation may be taken from an opposite side to inspire users from radical opposition and avoid hyper-personalization [Andersen and Knees, 2016]. Such an anti-solutionist [Blythe et al., 2016] approach to machine learning may encourage innovative developments that fundamentally reconsider the underlying notion of universal performance commonly at stake in the field of Machine Learning and arguably not adapted to the human users studied in the field of Human-Computer Interaction. It may also allow the building of imperfect AI colleagues, in opposion to "heroic" AI colleagues [d'Inverno and McCormack, 2015]: being impressed by the creative qualities of an abstract artificial entity may not be the best alternative to help people develop as creative thinkers [Resnick, 2007]. The *Co-Explorer* fairly leans toward such an unconventional design approach, which, in default of fitting every user, surely forms one of its distinctive characteristics.

Several machine learning frameworks remains to be investigated under the light of these human-centred challenges. Evolutionary computation methods [Fogel, 2006] may be fertile ground for supporting user exploration and automated refinement of example designs. Active learning methods [Settles, 2010] may enable communication flows between agents and users that go beyond positive or negative feedback. Dimensionality reduction methods for interactive visualization [Maaten and Hinton, 2008] may improve intelligibility of agent actions in large parameter spaces and allow for more trustable partnerships. Ultimately, combining reinforcement learning with supervised learning could offer users with the best of both worlds by supporting both example and feedback inputs. Inverse reinforcement learning [Abbeel and Ng, 2004] may stand as a technical framework supporting example input projection and transformation into reward functions in a parameter space.

5.8 Application I: *Riding the Co-Explorers*

We demonstrated the *Co-Explorer* at the *19th International Society for Music Information Retrieval Conference*, in Paris, France [Scurto et al., 2018b]. We designed an interactive exhibition stand in the form of a video game, that we called *"Riding the Co-Explorers"*.

5.8.1 Description

"Riding the Co-Explorers" is a single-player game that exists in a single-level form: it lets players guide the agent in one single sound space. The player's goal is to find a sound that they like. Thus, the video game possesses a free form: it does not aim at directing the player toward specific paths or locations in the level. This decision was specifically intended for the conference: we wanted our players to focus on the *Co-Explorer* more than on the video game itself, and thus uses the video game as a way to enhance immersion (see Video for an overview).

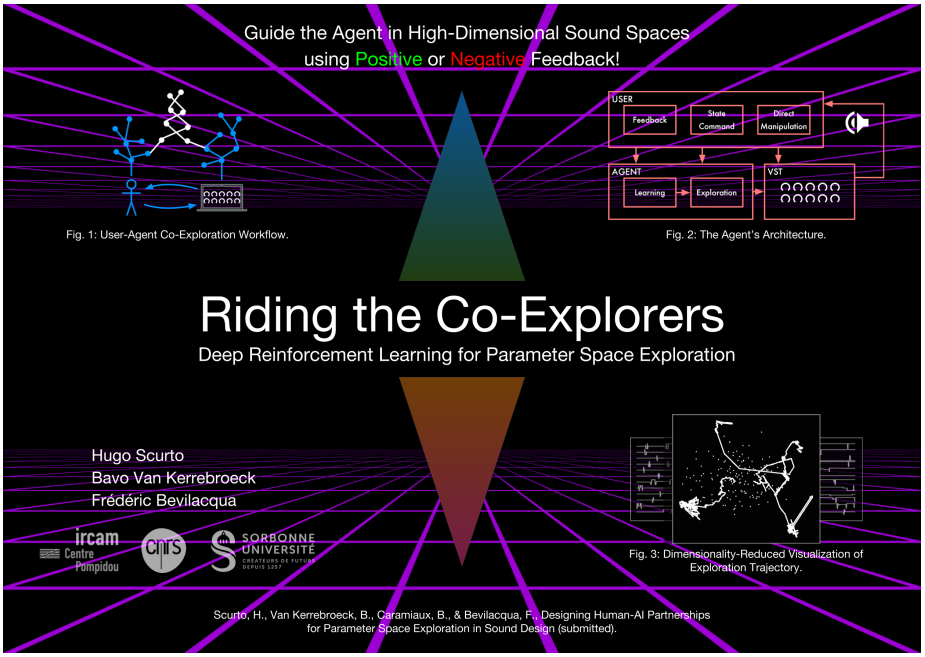


Figure 5.14: Visual theme for *Riding the Co-Explorers*.

The verb "Riding" aims at emphasizing the experiential aspect of interaction with the agent over the standard, task-related interaction paradigm that we tested in our previous study. In this sense, we reformulated the instruction given to users in standard interaction as a catchy storyline for the game: *"Guide the agent in high-dimensional sound spaces using positive or negative feedback!"*

"Riding", as it is used in horse riding, or motorcycle ride, also embodies the partnership that takes place between the human and the agent during sound space exploration. In this sense, we wanted to make the trajectory taken by the agent in

high-dimensional spaces visible to its human user. We opted for an abstract, three-dimensional representation of the agent and its environment, with aesthetic references to retrofuturistic video games (see Fig. 5.14). Specifically, we took the decision to let all past players' trajectories visible in the environment, to give a sense of diversity in the trajectories taken, and feedback given, by the players.

5.8.2 Game Design

Level

We selected ten parameters from a commercial VST to build a sonic space that would correspond to the agent's environment. We used the same parameterization of the *Co-Explorer* than we used in our previous study. As mentioned, agent memory consists in all players' trajectories and feedback given. This multi-user feedback context makes agent learning and behaviour much more complex than in the single-user use case previously investigated.

Graphics

We used Jitter to realize the graphics of the agent in its environment (see Fig. 5.15). We used Principal Component Analysis (PCA) to reduce the dimensionality of the agent's trajectory in the sound parameter space. We took the three first components of the PCA to project the agent in a virtual three-dimensional environment. We used the OSC protocol to send the PCA coordinates in real-time from the Python Script to the Jitter patch.

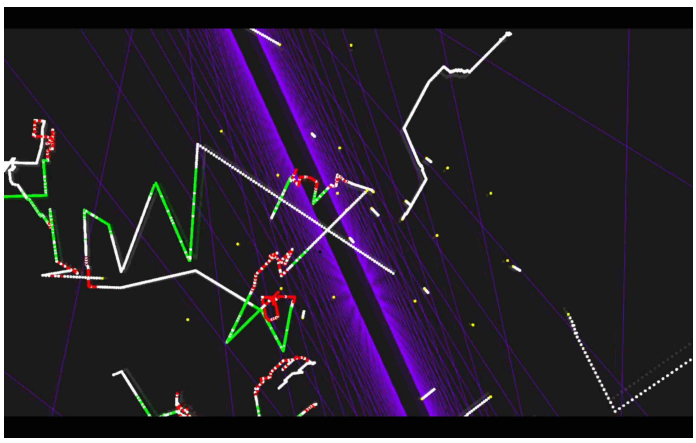


Figure 5.15: Screenshot taken from *Riding the Co-Explorers*.

User Interface

The user interface combines tangible and screen-based elements (see Fig. 5.16). An iPad allows players to communicate instructions to the agent. We used Mira to implement the iPad interface. Four buttons support guiding and zone feedback communi-

cation (red for negative, green for positive); a yellow button supports changing zone state command; a grey button allow players to reset agent memory; a white button enables players to start or stop autonomous exploration mode. Finally, a 20-inch screen enables real-time visualization of the agent's trajectory in the designed environment.

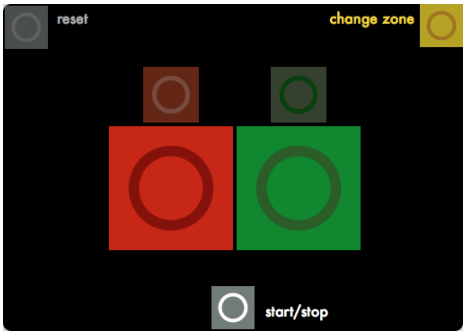


Figure 5.16: User interface implemented as an iPad-based tangible element.

5.8.3 Application in Exhibition

We exhibited *Riding the Co-Explorers* during three days at the ISMIR conference.

Stand

The exhibition stand consists in a minimalist space, typical of those used for demos in conferences (see Fig. 5.17). We put the user interface along with the screen and a pair of headphones on a table that was furnished by the conference. We printed the visual theme and hanged it behind the table to promote the game. The space required for the installation corresponds to approximately 3 meters long, 1 meter wide, and 2 meters high.

Attendees can spend as much time riding the *Co-Explorer* as they want to. They could use the installation with the help of a supervisor as well as without any supervision. When nobody was playing the game, we let the *Co-Explorer* behave in autonomous exploration mode to draw the attendee's attention toward our stand.

Reactions

Approximately eighty attendees played the game over the three days of the conference, trying to ride the *Co-Explorer*. The creative retrofuturistic visual theme worked well in catching the attention of the scientific and industrial attendees of the conference. Attendees spent between 1 and 5 minutes playing the game, ranging from standing random button pushing to seated immersed co-exploration.

As expected, our qualitative discussions with the attendees were more technology-oriented than game-oriented, which validates the one-level design of the game. Over the short period of time consacred to discussion, we were able to observe the same categories of users than during our previous study. Some players were user-as-leader



Figure 5.17: Picture of the exhibition stand used at *ISMIR'18*.

(skeptical about not always understanding what the agent does), some others were agent-as-leader (taking a ride on the *Co-Explorer* as it goes).

The visualization provoked mitigated reactions. Some attendees that were much focused on the sound produced by the agent took the visualization as a metaphorical representation of where the agent was in the space. Some others that were much focused on the technology in the agent were disrupted by the visualization, as they were trying to map the three-dimensional PCA position of the agent with the high-dimensional deep reinforcement learning architecture of the *Co-Explorer*.

We harvested quantitative interaction data along the three days of the conference. Unfortunately, it is hardly possible to extract any relevant information from it, as we were not able to witness agent learning along the three days of the conference. The reason is that most attendees needed to reset agent memory when starting interaction with the agent for the sake of clarity. Agent behaviour in a multi-user feedback interaction remains to be investigated more deeply, and could lead to interesting insight from both machine and human points of view.

5.9 Application II: *Behavioral Matter*

We applied the *Co-Explorer* to the design of robotic objects in a workshop called *Behavioral Matter*, held at Centre Pompidou, Paris, from 15th to 17th March 2019.

5.9.1 Description

Behavioral Matter, created and produced with the Centre Pompidou in the framework of the exhibition La Fabrique du Vivant (Mutations/Créations 3), was a multidisciplinary, international, and public research-creation workshop. It intended to question and test the notion of behavior in relation to matter and materials, objects and techniques, as well as to living or semi-living systems. The workshop stems from the "Behavioral Objects" project developed by the Reflective Interaction group of EnsadLab (the research lab of the École nationale supérieure des Arts Décoratifs - EnsAD, Université PSL, Paris) in partnership with the cluster "Matters of Activity. Image Space Material" from the Humboldt-Universität zu Berlin.

Over 70 artists, designers, researchers, student-researchers, students, and engineers from many different countries (Germany, Australia, Canada, Spain, the United States, France, Italy, Great Britain…) worked together during three days in the Forum, the main lobby of the Centre Pompidou (see Video for an overview). Twelve modules were designed to include all these actors, and to facilitate collaboration. With four other participants, I joined the module called Learning to Move, coordinated by Sofian Audry. The module explored machine learning applied to the development of the behaviors of objects. We were two doctoral researchers in the area of machine learning, one artist and teacher-researcher in the area of artificial intelligence, one teacher-researcher in the area of computational creativity, one teacher-researcher in the area of digital arts, and one student in the area of fine arts. My proposition was to apply the *Co-Explorer* outside the sound design domain by testing it on robot parameters to create movement behaviors.

5.9.2 Object Design

Process

The goal of the workshop was to test and open perspectives at the crossroads of several disciplines to develop works of art and design. The only requirement for the end of the workshop was to produce a public presentation of what was done in the module. We thus adopted an experimental and collaborative approach, and set the agenda collectively to coordinate our efforts.

The first day started by each actor making a short presentation of his or her work to the module, so that all members could better know each others with their expertise. We then familiarized with the MisB Kit[6], an open source building blocks kit that allows to quickly build and animate physical objects without requiring any previous experience in engineering, coding or robotics.

[6] http://misbkit.ensadlab.fr/



Figure 5.18: Picture[5] of the *Behavioral Matter* workshop at *Centre Pompidou*.

The second day was consacred to prototyping. We worked in teams of two on specific topics to tackle several points in parallel. Half of the module worked on supervised learning using the Wekinator[7]; the other half worked on reinforcement learning using the *Co-Explorer*. Moments were taken to give short tutorials to members that were not familiar to machine learning to try to balance expertise within the module.

The third day was consacred to documentation of the workshop and finalization of the created prototypes. We encountered wireless network issues with the robotic kit, which hindered communication between learning algorithms and robotic objects, and slowed the prototyping process a lot. We finally prepared a short speech and scenario for the public presentation, which took place in the end of the day.

Prototype object

We prototyped one robotic object using the *Co-Explorer* (see Fig. 5.19). It consists in a small, abstract object that learns to find its own equilibrium. The robotic thus exhibits a learning movement, rather than actually trying to learn to move.

It consists in one actuator, a servo motor, which can produce a rotational movement. We put one structural element, a shelf bracket, on the rotating surface of the actuator, using velcro. We fastened one flexible velcro bar to the structural element, on which we taped a piezoelectric sensor for flexion. Because of the gravitational field, the velcro bar will be subject to a bending. The goal of the robotic object is to learn which actuator angle would yield the less bending to the bar.

We used the *Co-Explorer* to learn the relationship between the actuator angle and

[7] http://www.wekinator.org/

[7] Taken from http://reflectiveinteraction.ensadlab.fr/.

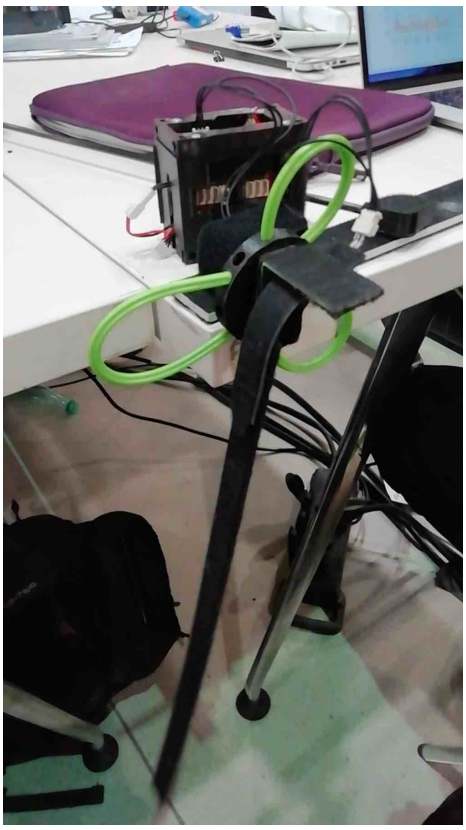


Figure 5.19: Picture of the robotic object.

the bar flexion. We modelled the angular position as the normalized state of the deep reinforcement learning agent ($S = s \in [0, 1]$). We defined the corresponding action space as making an angular movement up or down by one step $a_i = 0.01$, except when the state equals one boundary value.

We crafted a reward function by hand that is proportional to minus the square of the piezoelectric sensor ($R \propto -M^2$). The more the bar bends, the more the reward will be negative—the more the agent should change its angular position. We also tried a linear reward function ($R \propto -M$) that yields similar behaviour. We used the OSC protocol for communication between the *Co-Explorer* and the MisB Kit.

5.9.3 Application in Workshop

Results

Starting from either a bending or non-bending angular state, the robotic object manages to reach and converge to its goal state in a few minutes. The movement that we observed directly translates the reinforcement learning formalization that we adopted—*i.e.*, making small angular movements up or down. We thus built a self-learning object able to get a sense of its physics by means of computation.

We made another try after the robotic object learned to converge in its goal state. We held the velcro bar in a horizontal position to have the object believe that gravity now exerts its force horizontally rather than vertically. After tenths of seconds, the object left its goal state to explore other angular state, until finally converging to the new maximum-reward angular state. We also successfully tried to teach the object the good angular state by sending it human reward, as we did in the sound design application domain.

Insights

Even if quite limited from a robotic point of view, our results showed that the *Co-Explorer* supported adaptation to other parameter spaces than sound spaces. Its building on OSC protocol enabled rapid prototyping of small robotic objects using the MisB Kit. Its generic state-action parametric formalization supported learning in other spaces than sound, and from other reward signals programmable by hand or given from human trainers.

Future work may investigate whether the *Co-Explorer* yields relevant results on robotic objects that have more degrees of freedom—*e.g.*, possessing more than one actuator, and/or more than one sensor. One could also study dynamic state features instead of static angular state to observe different dynamic movements yielded by a reinforcement learning agent. Finally, working with non-expert users remains to be investigated more deeply, going beyond explaining its basis to fully get novice users work with it in an adapted pedagogical workshop structure.

5.10 Application III: *ægo*

We leveraged the *Co-Explorer* in a computer music piece, named *ægo*. The piece was designed through a research and creation process, in joint collaboration with Axel Chemla–Romeu-Santos. Our wish was to gather personal reflections from practice with machine learning to propose conceptual insight for future multidisciplinary inquiries in the realm of computer music. This work has been published and presented as paper and piece at the *14th International Symposium on Computer Music Multidisciplinary Research (CMMR 2019)* in Marseille, France [Scurto and Chemla, 2019].

5.10.1 Motivation

For this study, we were interested in adopting a *joint scientific and musical approach* to machine learning research. We are inspired by the computer music pioneer Jean-Claude Risset [Risset et al., 2007], 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 section 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.

5.10.2 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.

Unsupervised Learning for Sound Analysis and Synthesis

Musical Issue Most sound analysis-synthesis techniques, such as phase vocoder [Rodet et al., 1987] or wavelet transform [Kronland-Martinet, 1988], 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?

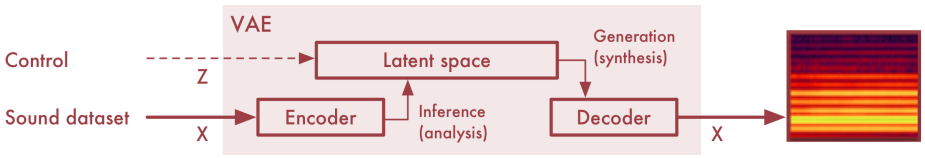


Figure 5.20: Unsupervised learning for sound analysis and synthesis. The variational auto-encoder (VAE) encodes a sound dataset into a high-dimensional latent space, which can be parametrically controlled to synthesize new sounds through a decoder.

Model The recent raise of *unsupervised generative models* can provide a new approach to sound analysis-synthesis, by considering each item of a given audio dataset $\{\mathbf{x}_n\}_{n\in 1...D}$, 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.

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 [Kingma and Welling, 2013]. In this model (Fig. 5.20), 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 criterions, such as enforcing perceptual constraints related to timbre [Esling et al., 2018]. We report the reader to the latter paper for technical details on the model and quantitative evaluation on standard sound spectrum datasets.

Reinforcement Learning for Sonic Exploration

Musical Issue Sonic exploration is a central task in music creation [Ystad et al., 2019]. 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?

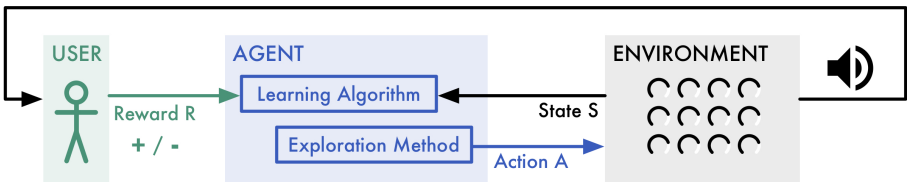


Figure 5.21: Reinforcement learning for sonic exploration. The agent learns which actions to take on a sound synthesis environment based on reward given by the musician. The agent implements an exploration method to foster discovery along interaction.

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 [Sutton and Barto, 2011]. 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.

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. 5.21). 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 [Warnell et al., 2017].

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 [Scurto et al., 2018a, Van Kerrebroeck, 2018] for technical details on the tool and qualitative evaluation from expert sound designers.

5.10.3 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.

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 [Assayag et al., 2006] 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 [Fiebrink and Caramiaux, 2016] that lets musicians experiment with the nonlinearities of the latent spaces.

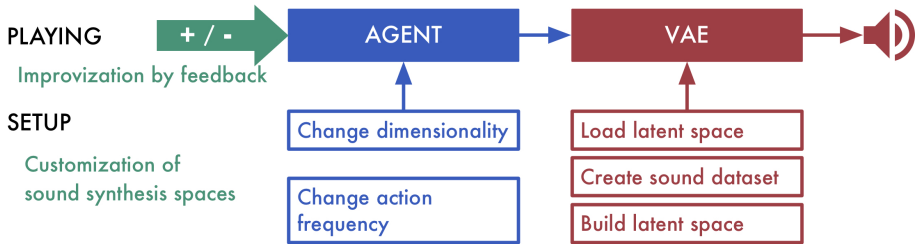


Figure 5.22: The interactive workflow that we designed for our instrument.

Workflow We designed a two-phase interactive workflow, shown in Fig. 5.22.

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.

Engineering

Implementation Technically (see Fig. 5.23), 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.

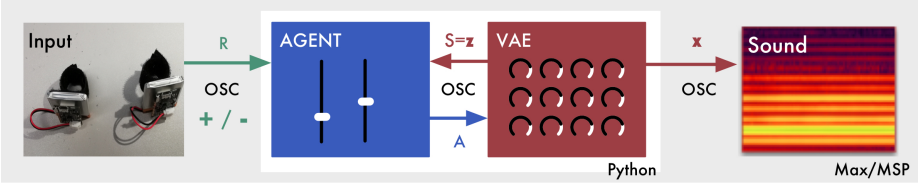


Figure 5.23: Graphical outline for the engineering of our instrument.

Hardware We designed a hardware prototype to support embodied musical interaction (see Fig. 5.23, left). It consists in two velcro rings, each of them equipped with

a wireless inertial measurement unit[8]. 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 [Tanaka and Donnarumma, 2019] to communicate positive or negative feedback to the agent.

Software We implemented our two machine learning models as Python libraries[9][10]. 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.

5.10.4 Musical Artwork

In this section, we present *ægo*, an improvisational piece that we wrote for our musical instrument (see Video for an overview). The piece was premiered at the *14th International Symposium on Computer Music Multidisciplinary Research (CMMR 2019)* in Marseille, France. We describe the intended aesthetics of sound, image and body, and detail how composition and performance were approached in relation to the learning abilities of our instrument.

Aesthetics

Motivation Our artistic motivation for the piece was to show and share an encounter between a learning machine and a human being with the audience. The learning machine possesses a latent sound space, as well as a distinctive musical behaviour, that are both originally unknown to the human being. Through improvisation, the human and the machine will learn to interact with each other—on an embodied level for the human, and on a computational level for the machine.

This mutual exploration is designed to be heard, seen, and experienced by the audience. The piece divides in two successive scenes (additive and physical), corresponding to two latent sound spaces learned by the machine. The performer will expressively negotiate control of these space with the machine, communicating positive or negative feedback using motion sensors placed in both hands. The slowly-evolving spectromorphologies, synthesized and projected in real-time on stage, intends to open a sensitive reflection on what is actually learned on a musical level, both by the human and its artificial alter ego—the machine.

Intentions The piece's aesthetic intentions toward machine learning lie at three intertwined levels: sound, image, and body (see Fig. 5.24).

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

[8] http://ismm.ircam.fr/riot/
[9] https://github.com/acids-ircam/variational-timbre
[10] https://github.com/Ircam-RnD/coexplorer

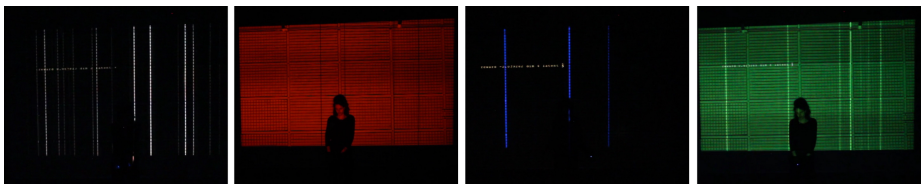


Figure 5.24: Pictures taken from the *ægo* artwork.

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.

Writing

Composition The piece was composed at three temporal scales (see Fig. 5.25).

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.

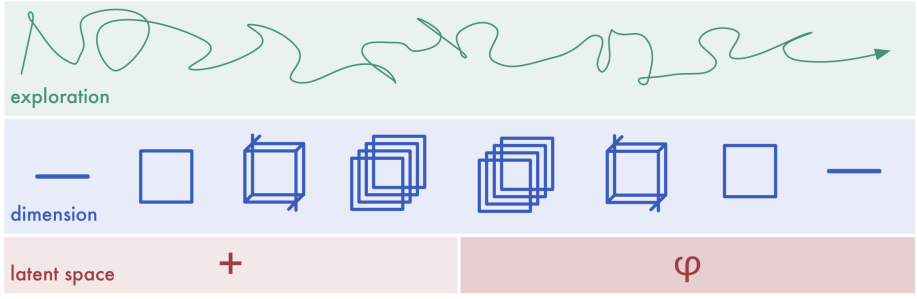


Figure 5.25: Temporal structure composed for the piece.

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 [Ballet et al., 1999]). 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. 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.

5.10.5 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.

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. 5.26, 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.

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. 5.26, dashed lines and arrows).

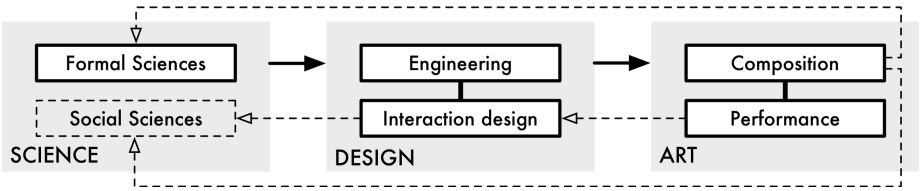


Figure 5.26: Our case study. Solid arrows: The sequential research and creation process that we took to scientifically investigate our models, and musically create our instrument and artwork. Dashed arrows: The personal conceptual insight gathered along our process.

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. 5.26, 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 [Chowning, 1973], 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. 5.26, 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 [Wiener, 1965]. 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 [Bevilacqua et al., 2009]. 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 [Risset and Wessel, 1999]. 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

Designing With Active Learning For Collective Musical Interaction

This chapter presents the fourth and last study led in the thesis. We decided to go beyond individual musical tasks studied in Chapters 3, 4, and 5 to investigate collective musical interaction. Our wish was to explore how machine learning could foster new types of collective musical interaction, possibly mediated by mobile technology.

We adopted a participatory design method to involve stakeholders at all stages of the design process [Schuler and Namioka, 1993]. Specifically, we started with an observation step to brainstorm interaction scenarios with stakeholders before deciding on the machine learning technique to be studied (see Fig. 6.1). This enabled us to frame active learning as a relevant technique for collective musical interaction. Yet, as we will see, only a reduced prototype was implemented in the context of this study.

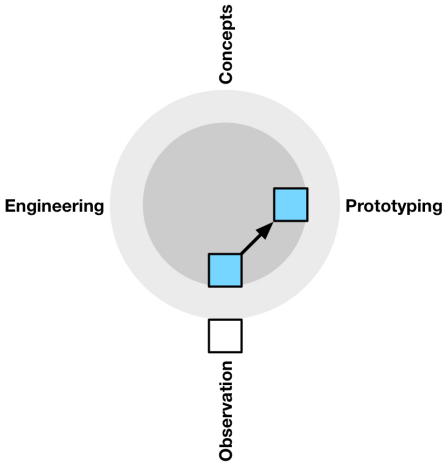


Figure 6.1: The participatory design method for our fourth study.

Section 6.1 reports on **observations** led in a field study, describing the musical task of collective musical interaction and interviewing designers of previous systems dedicated to this task. Section 6.2 describes the **design** of our model prototype. Finally, Section 6.3 describes the **application** of our model to the design of a public installation, called *entrain*, that we showcased at a five-day international conference.

The installation was published and presented as installation at the *ACM SIGGRAPH Studio 2019 (SIGGRAPH'19)* in Los Angeles, California [Scurto et al., 2019a].

6.1 Observation

We started our study with an observation step with designers of interactive music systems for collective musical interaction, especially, the *Coloop* system. Our wish was to focus on this system and brainstorm with our stakeholders how machine learning techniques could foster collective musical interaction.

6.1.1 Musical Task: Collective Musical Interaction

Before relating our study, we situate the musical task of collective musical interaction in the frame of contemporary music practices. We then present *Coloop*, the system on which our study will build on.

Situating the Musical Task

Collective musical interaction is a music task consisting in leveraging audio and visual cues to communicate expressive intentions to other people [Leman, 2008]. While a great skill of expert performers and music improvisers, collective musical interaction is also accessible to non-musicians [Leman, 2016]. For example, non-musicians may involve in collective music interaction when experimenting live music concerts. The expressive qualities of music enables musicians to communicate intention to non-musicians, and reciprocally. Yet, collective interaction is a complex human phenomenon, which remains hardly understood from a cognitive point of view.

In recent years, the *CoSiMa* project[1] adopted a design approach to investigate such collective music interactions by creating audiovisual experiences based on networking technologies. Crucially, one of the aims of *CoSiMa* was to leverage networking technologies—which were essentially addressed to expert composers and performers from computer music communities—to include a broader audience in music-making. In particular, the project build on web technologies to enable co-located audiovisual experiences based on mobile smartphones. Mobile technologies were shown to support the engagement of both musicians and non-musicians in collective musical interaction [Matuszewski et al., 2019].

As such, networked collective musical interaction may happen in public contexts, ranging form artistic installations to interactive experiences for urban services. It may relate to a different form of engagement than the three previously-investigated musical tasks, emphasizing expressive and experiential aspects of interaction over

[1] http://cosima.ircam.fr

expert creative usages. Our wish is to study how machine learning may help people engage in a music practice without necessarily wanting to engage in a use.

Coloop

Collective Loops [Schnell et al., 2017] and *Coloop*[2] are examples of these collective music-making systems in public spaces. *Coloop* allows up to eight persons to collaborate in the creation of rhythmic drum loops. To do so, each person's mobile phone is provided an instrument corresponding to a part of a drum set—*i.e.,* bass drum, snare drum, hihat, tom, or cymbal. Each phone's screen allows to fill a sequence of 16 beats with sounds corresponding to the instrument (see Fig. 6.2). The resulting rhythmic drum loop would thus consists in the addition of all eight persons' sequences synchronized in time. A 3D-printed loudspeaker equipped with LEDs renders the musical outcome in real-time, which enable participants to reassemble together in a shared place. Thus, *Coloop* intends to encourage participants to collaborate in the creation of danceable, or aesthetically-enjoyable rhythms.

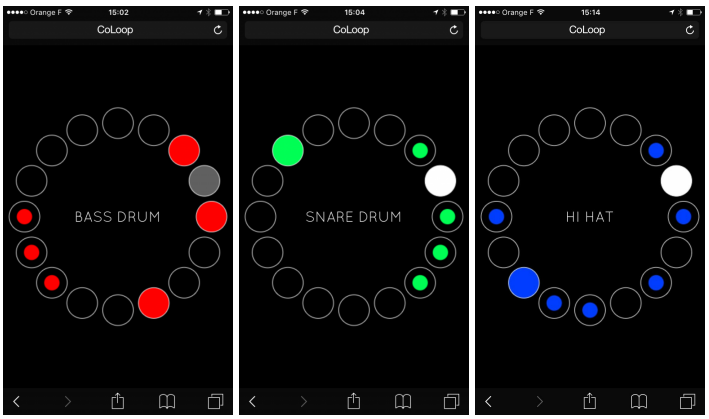


Figure 6.2: User interfaces for the *Coloop* mobile web application.

Coloop is a relevant use case to study collective musical interaction. First, from a human-centred point of view, it was one of the most successful audiovisual experience built during the *CoSiMa* project, as the installation was awarded a national design prize[3]. Second, from a machine point of view, its design remains quite generic in terms of implementation and musical affordances, which allows for easy incremental design actions. We thus decided to focus on *Coloop* to investigate machine learning applied to collective musical interaction.

6.1.2 Setup: Field Study

We led a field study with two participants that worked with the *Coloop* installation. The goal was to: Understand the strengths and weaknesses of the installation; Brain-

[2] https://www.nodesign.net/en/portfolio/coloop
[3] http://observeurdudesign2018.fr/labels/65-coloop-sequenceur-musical-collectif-connecte

storm ideas for possible design breawdowns.

The field study took the form of a joint semi-structured interview of the two participants. Our wish was to harvest personal reflections of the participants about the installation, while at the same time fostering collective discussion and having them exchange viewpoints. The first participant was a researcher and developer who actively contributed to the implementation of the installation. The second participant was a researcher and designer who had experience in supervising its functioning with the public. We logged the audio recording of the joint interview, who approximatively lasted one hour and a half.

6.1.3 Results: Qualitative Analysis

On *Coloop*

The first questions of the interview aim at asking participants about their previous experience observing the public interacting with *Coloop*.

Previous supervision. P1 and P2 respectively supervised the *Coloop* installation five and two times each. The public that engaged with the installation was essentially non-musician, young people. Contexts of interaction ranged from technology fairs to cultural evenings and art exhibits. Different sound designs were available within the installation; yet, the TR-808 and TR-909 designs were the most used by P1 and P2, as they were those that fostered the more interesting interactions.

The usefulness of a simple interface. When asked on the strengths of the installation, P1 and P2 jointly mentioned the simplicity of the smartphone-based individual interface. The non-musician public enjoyed creating rhythmic loops by only filling some of the sixteen blanks of the interface with sounds. As a matter of fact, three additional buttons that were implemented to generate automatic fillings (respectively, random, preset, or clear buttons) were almost never used by the public, who preferred filling their loop with their own beats. P2 said that this simplicity sometimes let non-musicians create *"very beautiful things. [...] It's not necessarily those that make techno music, that know what a drum machine is, etc. It's those that experiment with this very simple thing"*.

The lack of human learning. As a result of its simplicity, the installation did not really encourage the public to get a better understanding of its working. P1 evoked the fact that it is common place in public installations: *"Most people don't necessarily want to learn the system. So they may want to use it in its most straightforward version"*, P2 commented. As a consequence, there were many times where the eight public participants ended up filling all the blanks of the eight interface, which had the consequence of having one of the supervisor use a clear button (*"The most important button of the installation!"*, P1 joked) to restart the installation from scratch. While this filling could be analysed as a specific form of collective musical behaviour, one may argue that it digresses from *Coloop* original aim, which is making music collectively, without the help of a human supervisor.

On possible improvements

The rest of the interview progressively evolved in the form of a joint discussion to brainstorm possible algorithmic improvements for the installation.

Suggesting adapted loops within user interface. To encourage the public not to fill all the blanks, a first idea was to suggest specific fillings to users through their individual interface. While this idea may help guide the public to play different rhythms, P1 and P2 jointly agree that it would add too much complexity to *Coloop*, and have it lose its attractive simplicity. P1 inquired: *"It's less a matter of algorithmic complexity than of interface complexity. Imagine you being happy with the loop that you created, then something telling you 'hey I give you this'. How is the public supposed to listen to the suggested loop, while remembering its previously-created loop, and reflect on which one he would prefer?"*

Adding an adaptive, artificial player. To guide the public toward different rhythms, a second idea was to add an artificial player that would generate an additional loop that adapts to the eight loops played by the public. P1 noticed that such an adaptive behaviour would allow the installation to be usable by single persons, which could be helpful in cases where it would be exhibited in a hardly busy place. P2 underlined that this added adaptation would have the installation change its definition: *"I think that people must be informed about it, so that they may try to play with it, or even break it, or push it to play something. This would be better understood by the public [than suggestions], I think."*

Structuring the rhythm of interaction. A third idea was to structure the rhythm of interaction to clarify which of the artificial or human players would be responsible for generating loops. I suggested a simple question-and-answer format, where human players would play together during four measures to create some musical question, then listen to the artificial player during four measures generating an adapted musical answer—leveraging standard 4-measure structures at stake in electronic music. P1 and P2 disagree with this suggestion, saying that this would interrupt the public's appropriation of the installation, and that it may confuse non-musicians not aware of the 4-measure structure.

Reward the public with audio effects. A fourth idea was to reward the public with audio effects depending on some musical criteria determined by an adaptive agent. This rewarding would create a collective behaviour where all public participants would play in the form of a competitive game to be in possession of the audio effect. P1 and P2 both welcomed the idea of augmenting the actual human player's performance by an effect, instead of adding a supplemental rhythmic loop to the already rich musical output. Yet, P1 warned about the use of an objective criteria to rule musical interaction: *"Ethically, you are socially rewarded if you accept to be reified by a technological system. Well, if you present it as an artistic premise, why not..."*

Cultivating magic of the music installation. To counteract the biased concept of optimal performance applied to music, both participants gave the insight of cultivating magic in the music installation. P1 highlighted the interests of focusing on the interval between full artistic abstraction and full technical explanation for the non-musician public: *"We as musicians should accept the idea that the public may not understand everything, and that this is not a problem. This is the whole thing of interactive art. If you do not understand a thing, then it is boring. If you understand everything, then it becomes technical, which is not the same kind of object. In the interval lies magic, which can be exciting for the public."* P2 considered the possibility of using the technical adaptivity of machine learning to evoke a sense of ghostly presence of all previous public members.

6.2 Prototyping

Our observations enabled us to start designing a model prototype. In this section, we describe the active learning task, that we chose based on the design ideas suggested in our field study. We then present our model prototype, starting by the processing of human data to measure user performance within *Coloop*, and describing the *Bayesian information Gain* algorithm that was adapted to our active learning task. We finally describe its implementation as a web-based application.

6.2.1 Proposed Technique: Active Learning

We propose to frame the design opportunity in *Coloop* as an active learning task. In active learning, an algorithm sequentially query the user about the data it is processing [Settles, 2010]. These querying strategies have been shown to support more efficient learning in certain situations where human users possess expert knowledge on the task being learned. Our intuition was that these queryings may be adapted to the *Coloop* system to actively steer human users to make music collectively—*e.g.*, adding an artificial participant in addition to human participants. As previously mentioned, with respect to our participatory design method, we had not decided on the machine learning task to be investigated prior to interact with *Coloop* stakeholders.

6.2.2 Human Data: Performance

A first requirement for the active learning task to be realized is to define the data on which it may operate. Based on our field study, our wish was to build on the simple and successful interface of *Coloop*, allowing each human participant to control a 16-bit sequence (see Fig. 6.2). We thus designed two features aimed at extracting information on participants' performance with respect to our observations in our field study (see Fig. 6.3). Importantly, our approach remains general; one may use other features—and more than two—to describe a user's state in other collective use cases.

Our first design intention was to account for how participants may focus on their individual interface. We propose to measure *user activity* as the first feature of our model. User activity is based on the number of changes made by the user in the

sequence in each loop. A completely passive user would have null activity, while an overactive user would have maximum activity. The feature is computed at the end of each loop, which supports dynamic accounting of a user's state as he/she interacts with the system.

Our second design intention was to account for the filling of participants' sequences. We propose to measure *rhythmic periodicity* as the second feature of our model. Rhythmic periodicity is based on the autocorrelation[4] of the user's sequence. A fully-filled sequence would have minimum periodicity, while a large-spaced sequence would have maximum periodicity. Again, the feature is computed at the end of each loop, which supports dynamic accounting of a user's state as he/she interacts with the system.

6.2.3 Algorithm: *Bayesian Information Gain*

To prototype the active learning task, we used the framework of Bayesian Information Gain (BIG) [Liu et al., 2017]. Based on our field study, the general idea is to manage model adaptation to participants' behaviour.

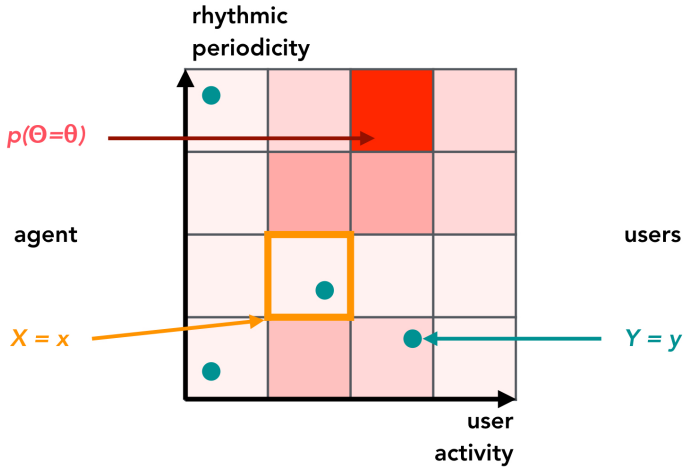


Figure 6.3: *Bayesian Information Gain* implementation. The two axes corresponds to user state features. The 4 dots represent the state of 4 different users. $P(\Theta)$, X, and Y correspond to the parameters in the BIG framework for adaptation.

BIG allows to adapt a probabilistic model of user behavior $p(\Theta = \theta)$ by actively sending some feedback $X = x$ to users and sensing participants' subsequent input $Y = y$ (see Fig. 6.3). The framework remains generic and thus potentially allows the machine to send any kind of feedback to participants—which may support the sending of audio effects. For our model prototype, we modified the framework to model the user state feature space. $p(\Theta = \theta)$ corresponds to the probability that a given *user activity* and *rhythmic periodicity* may correspond to the predicted user behavior. It is

[4] https://en.wikipedia.org/wiki/Autocorrelation

possible to define prior knowledge if we want the agent to have a positive bias toward a certain user behaviour—*e.g.,* medium activity with high rhythmic periodicity.

After some time of interaction, the agent may update its knowledge on users behaviour. The update is done by measuring users' move in the feature space, which corresponds to Y in the BIG framework. The agent updates its probabilistic model of user behavior and generates a new temporary feedback $X^{'}$. Importantly, the algorithm chooses X based on an information theory criteria. Specifically, it generates a feedback that may maximize the information gained from participants moves Y with respect to the probabilistic model $p(\Theta = \theta)$. We hypothesize that such a criteria may be hard to decipher for participants, and as such may cultivate some sense of magic within the installation.

6.2.4 Implementation: Web-based Application

We implemented our model prototype as a web-based application[5], building on the JavaScript implementation of *Coloop*. Human data processing and algorithmic computations are run on the server side. All mobile phones are registered as clients and share a unique clock, building on the *Collective Soundworks* framework for collective mobile web interaction. The resulting musical outcome—*i.e.,* the sum of all participants' rhythmic loops—exists as a client named "barrel". To start the application, one has to run a local server—*e.g.,* using NodeJS and a router. Then, all clients may be opened by accessing a web page—*e.g.,* using mobile devices for sequences, and a computer device as barrel.

6.3 Application: *entrain*

We applied our model prototype to the design of a musical installation, named *"entrain"*. We exhibited the installation at *ACM SIGGRAPH Studio* (SIGGRAPH 2019), in Los Angeles, California [Scurto et al., 2019a].

6.3.1 Description

entrain is a musical installation that lets participants make music together in collaboration with an adaptive agent (see Video for an overview). Participants may fill circular sequences to generate rhythmic loops—identically to the *Coloop* installation. Depending on their behaviour, an artificial agent may designate specific participants to steer collective music-making. *entrain* is inspired by the notion of *musical entrainment*, which refers to the human phenomenon of rhythmic synchronization that may occur when listening or playing music [Varni et al., 2011].

entrain combines active learning with hardware elements. Participants may use their own mobile phones to create their own rhythmic loops. A connected loudspeaker support sound generation, while also allowing participants to gather around a common physical place. A small single-board computer, placed inside the loudspeaker, support active learning from participants' behaviour to the adaptive agent.

[5] https://github.com/hugoscurto/entrain

The adaptive agent designates particular participants by sending audiovisual feedback in the form of audio effects and flashing lights. The audiovisual feedback may encourage social interaction between human participants and foster musical entrainment.

6.3.2 Installation Design

Hardware

For *entrain*, we collaborated with the design studio *Nodesign.net* to design hardware under the form of a connected loudspeaker and an embedded system (see Fig. 6.4).

Object. *Nodesign.net* designed a new loudspeaker, called *Coloop mini*. It reproduces the visual aspect of the previous *Coloop* loudspeaker with smaller dimensions. It was built using 3D-printing techniques. The upper part contains a white silicon patch, under which are placed eight LEDs.



Figure 6.4: Picture of the *Coloop mini* loudspeaker designed by *Nodesign.net*.

Embedded system. While the previous version relied on the bluetooth protocol to send audio information from a server to the loudspeaker, *Coloop mini* integrates a computer device at its basis, which is considered as an additionnal client to participants' mobile phones. A Raspberry Pi runs the web application as well as the "barrel" web page in real-time, and is directly connected to the loudspeaker with an audio jack. An Arduino Nano was programmed to have the barrel control the eight LEDs put over the loudspeaker.

Software

entrain leverages our model prototype to design the adaptive agent, by implementing specific musical structure and audiovisual effects.

Model update. We used a musical structure to update our model prototype in relation to agent's intervention. At the starting of interaction, the model prototype sets a feedback state X. The agent may update its knowledge on user behaviour every sixteen loops. The update is done by measuring participants' move in the feature space, which corresponds to Y in the BIG framework. The agent updates its probabilistic model of user optimal behavior $p(\Theta = \theta)$ and generates a new temporary target state $X^{'}$ that will be active for the next sixteen loops.

Audiovisual effects. The artificial agent intervenes in collective music-making by designating specific participants using audiovisual effects. We combined a delay with a pitch shifting effect to augment the designated participant's rhythmic loop. The delay was fixed at a ternary period to enhance the loop's musicality. We flashed the color of the designated participant using the LEDs of the loudspeaker, with rhythm corresponding to the participant's loop.

In the case where a participant gets into the feedback state during the sixteen-loop structure, the agent may make a "highlight" intervention by designating him or her with an audiovisual effect. This may act as a social incentive for participants: it encourages the designated participant to continue performing at the same level, while motivating other participants to play along with him/her to benefit from the same effect.

In the case where no user gets into the feedback state during these sixteen loops, the agent will make a "solo" intervention. First, it will stop interaction during four loops and designate the only participant that is closest to the temporary feedback state. This is aimed at switching participants' attention from their individual interfaces to the collective outcome. Second, it will slightly adjust user states toward the temporary target state. This may enable participants to start a new round with the installation, benefiting from additional motivation as provided by the artificial agent.

6.3.3 Application in Exhibition

We exhibited *entrain* at the *ACM SIGGRAPH Studio 2019*, in Los Angeles, California. Our aim was to harvest reactions on the interactive behaviour of the agent, understanding how its active learning may be perceived by different participants.

Setup

Participants. Approximately 500 participants interacted with the installation over the five full days or exhibition. As the conference's theme was Computer Graphics, few participants had experience or expertise with embodied interaction with sound. Actually, from our discussions and observations, apart from a few exceptions, most participants were non-musicians. We were thus interested in testing how non-musician people would react to our installation.

Configuration. Participants were presented with the installation as shown in Figure 6.5. To facilitate interaction, we provided them with mobile devices, that we displayed around the *Coloop mini*. To arouse participants' interest toward it, we made

the installation produce sound by pre-filling four mobile devices with loops. Then, the presentation setup varied depending on the number of participants that got involved at time of presentation. In most cases, participants came alone by the installation: either they had to play alone, or with us, or with unknown participants. Other participants came as a group to the installation and thus already knew each other before collaborating in music-making. When needed or asked, we introduced the functioning of the mobile devices' interface. Participants spent between one and five minutes interacting with the installation.

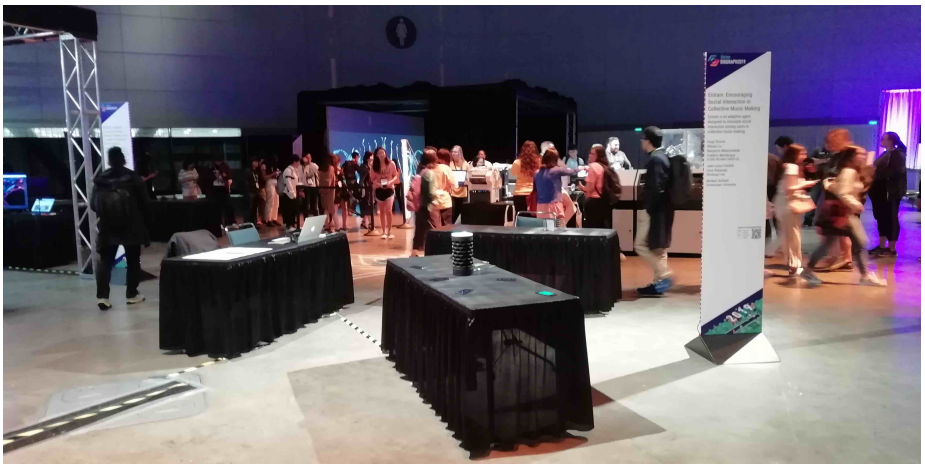


Figure 6.5: Picture of the *entrain* installation at *SIGGRAPH'19*.

Reactions

Participants in space. A first observation toward participants is the importance of spatial arrangement in relation to collective behaviours. For example, we left the mobile devices on the table to let participants use them. Hardly any participant ended up taking the device in the hand, which was our expected way of interacting with the installation. Also, once participants had chosen their mobile device, they stayed fixed to their initial position and did not move around the installation. This had us display the mobile devices all around the loudspeaker to encourage participants to enclose the loudspeaker and better perceive its sound and visual effects.

Participants' expectations. A second obsevation toward participants is the influence of their expectations in relation to their engagement with the installation. Participants' engagement at SIGGRAPH conference could be roughly classified in three levels. The first level corresponds to participants interested in understanding the technical functioning of the system. They typically asked about the model responsible for adaptation, as well as the web implementation and the realization of the loudspeaker, and did not really experience the installation as it was conceived for. The second level corresponds to participants interested in hands-on trying and playing with the

installation to try to make music out of it. Some of them were alone and constructed rhythms by gathering all mobile devices in front of them. Others collaborated with their colleagues or strangers to make music. The third and last level corresponds to participants in-between the two previous categories. They were not specifically interested in the technical specificities, nor the musical experiences of the installation, and thus had short experiences with the installation.

Presenters' influence over participants. A third observation toward participants relates to the influence of the presenters over participants. We noticed different participant behaviours depending on our interactions with the latters. For example, the simple fact of coming close to the installation when participants used it seemed to bias their behaviour. This may be explained by the fact that presenters may represent a form of authority in relation to the installation. Also, presenters' wording may completely change the reception of the installation by participants. For example, we began using the word "intelligent" to rapidly qualify the agent implemented in the installation. This seemed to hinder participants in experimenting with the system. Alternatively, using the phrasing "playing with you" to qualify the agent seemed to motivate our participants and foster their curiosity. We thus stuck to this phrase to complement the introduction of the installation.

Agent's functioning. Having all these factors in mind, we now report on participants' reactions toward the active learning agent.

First, the sixteen-loop update structure was rather good to create sequences of interaction that were not too short, and not too long. Participants had sufficient time to delve into individual sequence customization, while being stopped by the agent in case they did not enter the feedback state.

Second, the audiovisual effects triggered various reactions depending on participants' expectations. Most of the time, it fostered participants' curiosity toward the collective outcome of the installation. The "solo" interventions were appreciated by most participants, who imagined that the system reseted to allow a new series of collaborative music-making—in this sense, hardly any participant noticed that their sequence was being modified by the agent after this. Others that enjoyed the music outcome of the installation were entrained by it, in turn designating the participant designated by the agent. In case where participants were mostly focused on their individual interface, the "solo" intervention was perceived as a bug, as it stopped interaction with the screens. "Highlight" interventions were too quick to let participants identify which participant was being highlighted; as such, they were not looked after by participants, which were much more focused on creating their individual rhythmic sequence.

Third and last, agent adaptation generated different reactions among participants. First, no participant reported that they understood the adaptation mechanism in relation to collective music-making. This may validate our choice of information gain as machine criterion that may be hard for human participants to decipher. Second, active learning was successful in generating form during collective music-making. The algorithm often changed its feedback state after "highlight" interventions, which had

participants being designated not being designated anymore in a few seconds. Third, our model prototype converged quite quickly to some learned optimal user behaviour, which had us often reset its probability distribution by hand. This may pave the way for future improvements of the model in relation to collective human behaviours.

Discussion

Our application in exhibition showed that *entrain* created new collective musical interactions between essentially non-musician participants. The active learning implementation enabled to steer participants toward new musical configurations, while being sufficiently complex to appear as a black-box to them—which we believe was of interest for such an installation. The physicality of the loudspeaker helped participants situate themselves in collective interaction, as well as situating the agent in this configuration.

Future work may apply *entrain* in other contexts than that of conference to harvest situated feedback on the agent—*e.g.,* cultural events, club venues. Also, the model prototype may be improved to augment its robustness to user behaviour. Alternatively, other sound designs may be tested to generate variations in the collective musical outcome. For example, other audio effects may be used by the agent as a way to produce form during interaction.

7

On Human-Centred Machine Learning Applied to Music

In this chapter, we reflect on human-centred machine learning applied to music by analysing the four studies led during the thesis. We first analyse the harnessing of human musical goals and values that a human-centred approach to machine learning brings. We then assess the switching of machine learning tasks for general interaction techniques, as enabled by our introduction of model prototypes. We finally introduce the notions of *machine expression* and *co-expression* to examine the human-machine workflows designed with machine learning along our studies.

7.1 Harnessing Human Musical Goals and Values

The human-centred approach to machine learning allowed us to harness human musical goals in situated musical practices, as well as values considering the sociotechnical aspects of learning algorithms.

7.1.1 Human Feedback on Musical Tasks

The first advantage lies in the harnessing of situated human feedback in the design process. By situated, we mean feedback harvested in different human contexts, relating to (1) existing musical practices, (2) degrees of musical expertise, and (3) designer.

Feedback from Expert Music Practitioners

Designing with musical practitioners helped us include useful feedback in the model design process. Music practitioners are an extreme category of users [Garcia, 2014]: they accumulate hours of practice related to their art, which provides them with deep and specific expertise on it.

Deep, because they may know general dimensions that may be relevant to the practice of specific elements. For example, the expert movement practitioners with which we worked in the design of unsupervised learning for motion-sound mapping

(Chapter 3) leveraged their embodied knowledge to assess a novel interactive situation linking their gestures to sound. Including this feedback may provide high-level insight on a given musical practice, which in turn helps to generate ideas for the interactive affordances of a model prototype.

Specific, because they may exhibit a unique style in the execution of a given task. For example, the sound designers with which we collaborated in the design of deep reinforcement learning for synthesis exploration (Chapter 5) all had specific strategies for exploring sound spaces with standard parametric interfaces. Including this feedback provides low-level insight on the various steps of execution of a given musical task—for example, the spontaneous exploration of parameters—, which in turn helps to explore various high-level features of a model prototype—for example, the exploration strategy of the deep reinforcement learning agent.

Feedback from Non-Musicians

Perhaps counter-intuitively, non-musicians may be extremely useful to provide alternative feedback in the model design process. Even if they are not practitioners, non-musicians may experience music on a more or less regular basis, for example during listening, or in live venues [Leman, 2008], which provides them with general and diverse views on it.

General, because they may provide feedback related to more common issues related to music, rather than on practice. For example, the participants that we observed interacting with reinforcement learning for sonic exploration (Chapter 4) enabled us to gain an understanding of how humans may perceive agent exploration. Including this feedback may enable the selection of given alternatives in the design space, for example by excluding features that are not perceivable by humans.

Diverse, because they may witness novel feedback that may not originate from the state-of-the-art, nor from expert practitioners. For example, the non-musicians that we observed testing active learning for collective musical interaction (Chapter 6) witnessed a higher sensibility to contextual factors—such as system's expectation or presenters' influence—than expert musicians—who may have extensive experience with testing music systems.

Feedback from Designers

Designers of interactive music systems, as diverse as they may be (as described in Section 2.1.2) may provide valuable feedback in the design of a model prototype. This might seem obvious at first glance, as designers often implicitly discuss with, of sometimes belong to, the research and development group that carry the work. Yet, explicitly including them in the design process may benefit the process itself, but also the research community in its entirety.

The process, because designers may provide feedback that encompass conceptual and technical dimensions related to an element of a music practice. For example, the designers with which we discussed to design active learning for collective musical interaction (Chapter 6) were able to identify the key conceptual issues related to the

building of collective music patterns between participants, along with the technical challenges related to measuring participants' behaviour.

The research community, because most of the time, designers may be in the position of the observer, not of the observed. As a consequence, this feedback may not always be transcribed in most scientific publications, which prevents other researchers and designers to benefit from it. For example, the designers with which we discussed to design active learning for collective musical interaction (Chapter 6) were able to elicit the actual design flaws at stake in their own interactive music systems, perhaps a bit more than in scholarly publications.

7.1.2 Music Practices and Machine Learning

The second advantage lies in the sociotechnical consideration of machine learning in relation to music practices. By sociotechnical, we mean three aspects: (1) the relevance of automation in relation to expert music practitioners, (1) the biases that humans may have toward technology and "artificial intelligence", and (3) sustainable music practices.

Automation and Expert Music Practitioners

Machine learning has the potential automate several musical tasks (as we saw in Section 2.2). Yet, automating musical tasks cannot be considered without considering expert music practitioners that actually carry these tasks. The latters' practices may have conceptual and cultural specificities when apprehending automation.

Expert music practitioners have long expressed different opinions on automation. On the one hand, some expressed the wish to create machines that automatically create music [Hiller Jr and Isaacson, 1958], which is nowadays perpetuated in fields such as Computational Creativity, or even Musical Metacreation [Pasquier et al., 2016]. On the other hand, perhaps by a fear of automation in relation to composition [Dubnov and Surges, 2014], other composers expressed the wish to rely on rule-based programming to build machines able to assist the act of composition, such as with the OpenMusic software [Bresson et al., 2011]. This polarization was witnessed in our study of deep reinforcement learning for synthesis exploration (Chapter 5), where professional composers led creative tasks by positioning between an agent-as-leader—corresponding to automation-oriented composers—and a user-as-leader—corresponding to rule-oriented composers—workflow.

Having this conceptual framing in mind, cultural factors also have to be taken into account. Expert music practitioners often rely on embodied knowledge along their creative processes, leveraging specific strategies and values to explore the space of possibilities [Andersen and Knees, 2016]. It is important to stress that this embodied knowledge may be highly shaped by cultural elements—such as historical background on computer music—, and as such, may not be shared by all expert music practitioners. Automation in machine learning should thus be designed with care to account as much as possible for a diversity of cultural processes, which are the signature of human creativity. For example, our study of deep reinforcement learning for synthesis exploration (Chapter 5) showed that automation was appropriated by human users to

provide flexible paths to explore the design space, as opposed to simply replace them with one single automated path.

"Artificial Intelligence Technology" and People

The idea of automation may also be perceived differently by non-musician people, which may not have the same background than musicians when approaching technology and "artificial intelligence technology" in relation to music.

People may have biases on the notion of technology. Some of them may perceive technical objects from an ontological point of view, that is, objects that extend their natural abilities [Simondon, 1958]. Others may express fear toward technology, based on the belief that it has the power of changing the world [Haynes, 2003]. These cultural distinctions in perception continue to take place with the renewed rise of "artificial intelligence technology". The notion, being spread though science-fiction centuries ago [Buchanan, 2005], and investigated by scientists around half a century ago [Minsky, 1974], perhaps strengthen the division between technology-suspicious people, and technology believers [Geraci, 2012]. In both cases, this may have consequences on the perception of new technology by human users. For example, our study of deep reinforcement learning for synthesis exploration (Chapter 5) showed that our participants perceived the Co-Explorer as being not that "intelligent", which may suggest that they have important expectations toward the behaviour of the system. These expectations should be taken into account when designing technology that use machine learning.

Despite these biases in perception, some aspects of interaction may be shared by a greater majority of people, even non-musicians, notably in relation to creativity and flow [Csikszentmihalyi, 1997]. The music cognition field offers a great description of how expressive dynamic processes and feedback loops may actively arouse motivation and flow states in musicians [Leman, 2016]. In this sense, automation may be used to design a sense of play to foster human creativity instead of focusing on productiveness. This consideration backs the notion of improvization, as it was described in Section 2.2.2. As an example, the somasticks, designed from our study of unsupervised learning for motion-sound mapping (Chapter 3), provided with a degree of automated interactivity that seemed to be enjoyed by creative practitioners.

The Challenge of Sustainable Music Practices

A third and following sociotechnical aspect relates to the sustainability of the created music practices. Here, "sustainable" is meant to evoke the longer-term acceptance of machine learning technology by musicians and non-musicians. While not investigated in the context of this thesis, we stress the importance to lead longitudinal studies to assess the impact of machine learning on both mainstream and niche musical communities.

Mainstream, because machine learning has received considerable attention from the academic communities as well as the music industry. While these communities both consider the creation of new musical uses, they may have a tendency to do so through the prism of human productivity, whether it be at an individual scale—*e.g.,*

creativity support tools for music production [Deruty, 2016]—, or at a larger scale—*e.g.,* recommender systems for mass music consumption [Montecchio et al., 2019]. While partly motivated by scientific evidence, one can suggest that this drawing of attention may also be encouraged by sociopolitical pressures toward applying artificial intelligence [Caramiaux et al., 2019]. Research may be essential to study machine learning at the scale of mainstream music production, to identify which of the new musical practices may remain sustainable in such a highly-disrupted context.

Niche, because the attention of machine learning for mainstream music communities may harm smaller niche music communities. From a sociological point of view, music practices always inscribe within a cultural context, which manifests itself through shared knowledge and expectations on the sonic aesthetics [McLeod, 2001], as well as shared know-how and appraisal on compositional and performative techniques [Frith, 1998]. To sustain the diversity of music practices implies accounting for a diversity of human users in the design of machine learning. Data-driven techniques investigated in this thesis may go in this sense, but remains to be investigated on a longer term to demonstrate their impact.

7.2 Switching Machine Learning Tasks for Interaction Techniques

The human-centred approach to machine learning allowed us to switch machine learning tasks for general interaction techniques to be designed [Beaudouin-Lafon, 2004]. We detail how model prototypes helped such switching by generating new design ideas, and how the reassignment of model engineering in the design process allowed the concretization of these ideas.

7.2.1 Model Prototypes as Design Artifacts

The first advantage lies in the use of model prototypes as actual artifacts of the design process [Beaudouin-Lafon and Mackay, 2009]. This enables to frame new computations at three levels: (1) exploring new machine learning tasks, (2) expanding data types, and (3) generating new algorithms.

Exploring Open-Ended Music Tasks in Machine Learning

Leveraging model prototypes helped us explore new machine learning tasks in relation to their musical application. Machine learning models are often aimed at realizing a goal-oriented task with the highest performance in a given software application. In addition to these goal-oriented tasks, prototypes enable to explore open-ended tasks, which can be of interest for the creative domain of music.

Goal-oriented tasks are the main point of focus in machine learning literature, which developed optimization algorithms to add to the basic statistical models aimed at identifying patterns in raw data. The decision that the model should take from these identified patterns—*i.e.,* the machine learning task—remains quite flexible to design, computationally speaking [Bishop, 2006]. Most of the time, machine learning

tasks are goal-oriented as a logical continuation of their statistical modelling origin: examples are classification or recognition. This type of task can be useful in a musical application: for example, our study of reinforcement learning for sonic exploration (Chapter 4) showed that an agent having the goal-oriented task of maximizing feedback helped human users to effectively attain a musical goal. As they constitute the majority of machine learning literature, experimentation outside these remain quite rare.

Open-ended tasks, on the contrary, constitute a minority of the machine learning literature. They consist in taking decisions from patterns identified in data that go beyond optimizing some given measure to potentially discover novel patterns or generate novel behaviors. Using model prototypes allows to explore such open-ended tasks in the early stages of the design process, by assuming the possibly non-optimal behaviour of the model. For example, our model prototype used in our study of active learning for collective musical interaction (Chapter 6) allowed us to consider the open-ended task of stimulating dynamic music processes among human players. This kind of task may be promising for music, as well as for any other creative application domain, which, as discussed above, may often rely on dynamic, open-ended processes, than on static, goal-oriented outcomes.

Scaling Models and Data to Music

Using model prototypes also enable to scale models and data to musical applications. Most of current applications in the machine learning literature focus on using deep models on well-established data bases. By looking at model prototypes, we may investigate shallow models, while also witnessing new data modelling opportunities.

Model architecture is central to the development of machine learning. Specifically, most recent breakthroughs of artificial intelligence have built on complex, deep models trained on huge databases [Briot et al., 2017]. Alternatively, model prototyping encourages investigation of shallow models able to learn from small scales of data. These small scales are of main interest for music, where customization is of prime interest for creators [Fiebrink and Caramiaux, 2016]. The unsupervised learning model that we designed for motion-sound mapping (Chapter 3) may be the most representative example of such shallow model, as it leveraged a few seconds of human motion observation to create expressive musical interactions in the somasticks. Reciprocally, model prototypes can also be used to hijack deep models with a small data set approach, as we did in our study of deep reinforcement learning for synthesis exploration (Chapter 5).

Data at stake in machine learning can pertain to different signal categories. Some disciplines, such as Music Information Retrieval, may focus on a given categories of data—*e.g.,* data from audio signals—to do machine learning research. While it may sound like an obvious statement for other disciplines, we recall that many machine learning applications could be discovered if one considers new types of data and investigate them. Discovering and studying such new types of data could be an opportunity for music applications of machine learning to go from lab-based setups to real-world domains [Wagstaff, 2012]. For example, using a model prototype in our study of reinforcement learning for sonic exploration (Chapter 4) allowed us to con-

sider non-stationary reward signals through human feedback data, which is barely studied in the reinforcement learning literature.

Balancing Model Performance with Musical Features

Model prototypes enable to rapidly experiment interaction between machine learning systems and humans. As such, they enable to better balance model performance with the musical behaviour provided to humans in music activities. While the latter is the specificity of crafting in interactive music systems [Jorda, 2005], the former constitutes most of the focus of machine learning research.

Performance is how well a machine learning model performs a given task after being trained on a given data set. It is one of the main criteria for assessing the relevance of a model on a given machine learning task [Murphy, 2012]. Yet, this notion may not always be relevant for applications in music. The notion of interactivity was introduced in interactive music systems to go beyond the reactivity of standard, task-oriented interfaces [Paine, 2002]. One could think of a similar analogy in interactive machine learning applied to music. Our study of deep reinforcement learning for synthesis exploration (Chapter 5) showed that an artificial agent may be useful for a musical task even if not performing the learning task one hundred percent accurately. Performance also links with the notion automation in relation to music practices, as discussed in Section 7.1.2.

In addition to interactivity, explainability may be a relevant notion to the design of machine learning for music. Explainability is how interpretable the decisions a of a machine learning algorithm are for humans [Ramos et al., 2019]. It may be crucial to build trust in humans involved with technology in creative activities. For example, our study of deep reinforcement learning for synthesis exploration (Chapter 5) showed that relying on interactive visualizations—in this case, interface parameters and navigation history—allowed humans users to better understand the behaviour of the algorithm, and to better exert human control over algorithm performance. The use of model prototypes enabled us to switch from a performance-oriented to a human-trustable model design process.

7.2.2 Reassigning Model Engineering In The Design Process

The second advantage of a human-centred approach lies in the reassigning of engineering toward the end of the music design process, as opposed to standard engineering of machine learning (see Section 2.2.3). This enables to: (1) move from prototype to final models, (2) better seize engineering in relation to musical behaviour, and (3) move toward machine learning design patterns.

Centering Model Engineering On Humans

The human-centered approach enables to center the engineering machine learning on human musical practices. Both training and tuning of the models are done by including human users in the loop.

Training consists in collecting and labelling data sets from which the machine learning algorithm would learn a model. Often, these data sets are constituted by the researchers or designers of machine learning systems themselves, with the aim of providing a representative sample of the phenomenon to be modelled. As a consequence, the resulting big data sets are often impersonal, and potentially reproduce biases already present in the real-world. The human-centered approach allows to explicitly take into account the training process by directly having end-users build their own data sets interactively. For example, our study of deep reinforcement learning for synthesis exploration (Chapter 5) showed human users provide feedback labels to train different agents for different creative tasks. Customization may be essential in musical applications, as each musician may exhibit different creative processes and, as such, be considered as a unique phenomenon.

Tuning consists in adjusting the hyperparameters of a machine learning model before the training process. Often, it is done with the aim of optimizing the model behaviour in relation to the loss function of a given goal-oriented task. Yet, as discussed above, open-ended tasks, where a loss function may not be easily defined, may also be of interest for musical applications. The human-centered approach enables to explicitly consider the tuning process, either by choosing hyperparameter values from feedback-driven heuristics, or by directly having end-users choose their values. For example, our design of unsupervised learning for motion-sound mapping (Chapter 3) let human users choose the number of Gaussian components of their machine learning model, as well as the temporal duration corresponding to the batch size. Altogether, training and tuning allows end-users to appropriate the engineering of machine learning and use it for their creative purposes.

Model Engineering and Musical Behaviour

Despite our focusing on model prototyping, it is important to recall that model engineering still has consequences on the musical behaviour of a model. To analyse these consequences, one may distinguish between end-user and feature engineering.

End-user engineering refers to the interactive training and/or tuning of the machine learning model by the end-users of a system. It is the approach that we decided to focus on in this thesis, inspired by the field of interactive machine learning [Amershi et al., 2014]. End-user engineering may be central to the musical behaviour of a model. By allowing humans to customize a model and explore the space of possibilities, end-user engineering may be essential for humans' creative process in open-ended tasks. For example, our design of deep reinforcement learning for synthesis exploration (Chapter 5) let human users experiment with various training and tuning configurations to explore and discover new sound designs.

Feature engineering refers to the definitive training and/or tuning of the machine learning model by developers of a system. It is the step that we decided to jump over in our research methodology, except in our study of deep reinforcement learning for synthesis exploration (Chapter 5). Yet, feature engineering importantly influences the musical behaviour of a model. By providing humans with efficient interaction techniques and automation, feature engineering may be crucial to improve human productivity in goal-oriented tasks. Future work may push the design of our four

machine learning models further by leading feature engineering based on the observations harvested, while keeping in mind the sociotechnical factors related to machine learning and productivity raised in Section 7.1.2.

Toward Music Design Processes for Machine Learning

To frame new machine learnings for music requires considering both end-user and feature engineering processes in joint relation with situated human feedback and real-world applications. Design processes that explicitly take into account the time needed for conceptual and technical framing should be conceived.

Conceptual framing supports the formulation of a music practice as a machine learning task. This formulation includes defining a machine learning task, choosing a learning algorithm, and delimiting human data. Many possible strategies could be conceived for such reformulations. Theory provides state-of-the-art considerations toward the musical practices, and also the use of certain machine learning techniques. The human-centred approach integrates human users in the design process by having them discuss and test potential formulation in the form of model prototypes. Alternatively, our study of active learning for collective musical interaction (Chapter 6) involved human users in all aspects of conceptual framing. In the latter, human users were able to understand and formulate machine learning tasks in relation with the musical practice—even if not experts of machine learning. This suggests that human users can be included from the beginning of the design process for machine learning.

Technical framing enables to assess the working of the machine learning system. This assessment includes the testing of end-users engineering in interactive setups, as well as the feature engineering of the machine learning model in non-interactive setups. While the former only requires programming to be effective, the latter often takes a lot of time due to the big data needed to train and tune the machine learning model. These big data scales often slows the design process of the model, as constituting such big databases is actually a research contribution in itself. Using model prototypes enables to reduce this time by directly evaluating the learning requirements of the machine learning model in relation to the musical practice. Alternatively, teamwork is a reliable option to make conceptual and technical framing parallel. Our study of deep reinforcement learning for synthesis exploration (Chapter 5) stands as an example of such task assignment, having one researcher lead conceptual framing and another leading technical framing.

7.3 Creating Co-Expressive Human-Machine Practices

Having discussed the advantages of human-centred machine learning applied to music, we now analyse the musical interactions that it created. To do so, we propose to adapt Leman's notion of expression—which was essentially developed to describe interaction between humans and music [Leman, 2016]—to the case of human-machine interaction—specifically, where the machine learns to produce musical output from interaction with humans. We introduce the notion of *machine expression* to account

for the perceived musical behaviour of learning machines in relation to humans. We then analyse *co-expression*, as the dynamic process of musical expression between learning machines and humans. We finally leverage these notions to discuss artificial creativity from the point of view of human expression and music.

7.3.1 Machine Expression

The notion of expression is not new in the field of computer music. Some researchers even focused on expressiveness as an essential feature of interactive music systems [Jorda, 2005], notably through the spreading of NIMEs—acronym for *New Interfaces for Musical Expression* [Dobrian and Koppelman, 2006]. Yet, we believe that the "E" in NIME has mainly been considered for members of its computer music community—*e.g.,* expressiveness for expert performers, or for expert audience members. We are interested in including more people in the notion of musical expression, especially through the design of interactive music systems.

The notion of *machine expression* aims at describing interaction with interactive music systems from a human-centred point of view—be they musicians, or non-musicians. Machine expression pragmatically addresses the fact that humans may perceive expression in machines, regardless of machines' abilities to express by themselves. As such, machine expression may help support a wider range of human creative processes in the design of interactive music systems. In this section, we (1) situate machine expression in standard interactive music systems, describing the expressive affordances of computer devices and music itself as perceived by humans. We then (2) analyse machine expression in the context of machine learning, where the notion of artificial creativity has driven much of the research on interactive music systems. We argue that the notion of machine expression enables to design more inclusive machine learning for human creativity support in music.

Situating Machine Expression In Music Dispositifs

Marc Leman defines expression as "quality of human movement that is meant to provoke an expressive response, so that an expressive interaction can be established between sender and receiver" [Leman, 2016]. Here, we explain how machines may be perceived as senders of expressive signals by human receivers, considering the embodied qualities of computer devices and music in relation to human sensory-motor state (as depicted in Fig. 7.1).



Figure 7.1: The embodied qualities of music devices support the sending of expressive signals from machines to humans (as will be detailed in Figures 7.2 and 7.3).

Computer devices may be described as possessing embodied qualities related to their dynamic attributes and affordances. On the one hand, dynamic attributes refer to the exhibition of a movement behaviour from a computer device. It may apply to software in the case where the user interface exhibit autonomous dynamic behaviour. For example, our design of deep reinforcement learning for synthesis exploration (Chapter 5) expressed dynamic attributes in the user interface in the form of an agent moving parameters autonomously. It may also apply to hardware if the device is capable of robotic behaviour. On the other hand, affordances are what both software and hardware elements may offer to the individual [Gibson, 1977, Gaver, 1991]. They have been greatly investigated in computer devices through the wave of embodied interaction [Dourish, 2004]. For example, our study of active learning for collective motion-sound mapping (Chapter 6) leveraged smartphone devices that jointly express interface- and object-based affordances.

Music may be described as possessing embodied qualities related to its sonic forms and affordances [Caramiaux et al., 2014b]. On the one hand, sonic forms refer to the spectral qualities of sound organized in time—*e.g.,* rhythm, melody, timbre, or harmony. They have been shown to be associated with expressive gestures by humans, notably through the theory of embodied music cognition [Leman, 2008]. For example, our study of unsupervised learning for motion-sound mapping (Chapter 3) shows how drumming sounds express embodied qualities related to human movements to human players. On the other hand, affordances refer to the ability of music to trigger specific movements in humans. They have been used to align the rhythm of human movement on the music rhythm, notably through the notion of entrainment [Leman, 2016]. For example, our study of active learning for collective motion-sound mapping (Chapter 6) used rhythmic loops to express accessible musical affordances to human players. Other works in the field of sonic interaction design extended the notion of music affordance to that of sonic affordance, that is, affordance exhibited by sound itself, regardless of musical qualities [Altavilla et al., 2013]. We will only mention that one may consider either sound or music affordances in machine expression, depending on their field of work.

Analysing Machine Expression From Human Perspectives

As a general notion, machine expression may be analysed in a wide variety of existing interactive music systems (including the reactive tools, partners, and mediums described in Section 2.1.3). In the context of this thesis, our interest lies in analysing machine expression from a human-centred perspective. That is, our wish is to account for humans' perception of expression in machines regardless of their musical expertise. Such an analysis may enable to include more people in the design of interactive music systems. Here, we leverage the biosocial and cultural dimensions of expression [Leman, 2016] to analyse how machines may be perceived as expressive by humans—be they musicians or non-musicians. We depict the resulting concepts in Figure 7.2.

Biosocial signals originate from sensorimotor reflexes, triggered by sound perception, affecting the cognitive state of humans and animals. The most basic example of biosocial signal may be described through the notion of *entrainment.* Entrainment is

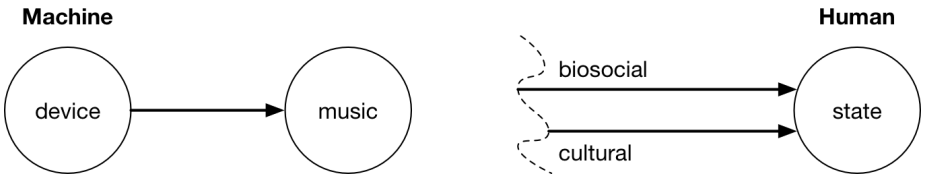


Figure 7.2: Machine expression considers the biosocial and cultural dimensions of music dispositifs to include diverse human perspectives on their perceived behaviour.

the process of rhythmic coordination that can take place between a living organism and an external stimulus [Phillips-Silver et al., 2010]. In the case of a sonic stimulus, coordination between the organism's movement and sound may take place, as a byproduct of pre-existing biological adaptations. In the case where the stimulus originate from another entity, coordination between the organism's movement and the entity may take place, as a byproduct of embodied associations and social feedback loops. Importantly, biosocial signals may not necessarily require the sender to have the intention of conveying an information; rather, they may be only based on the receiver's dispositions to sensorimotor responses. As such, machines may be analysed as emitting biosocial signals from the perspective of humans. For example, our study of deep reinforcement learning for synthesis exploration (Chapter 5) showed that humans made biomorphic associations toward the agent's actions on the VST.

Cultural signals involve some form of human learning toward biosocial expressive signals [Leman, 2016]. They may relate to the evolution of human cognition toward sound, for example regarding its intentional use as a communication modality, especially through the invention of language, of which music may be thought as being part of [Levitin, 2006]. Cultural signals may also relate to the development of human norms toward music, for example regarding the assessment of its value, which heavily depend on its situation in the communities to which it pertains [Small, 1998]. Machines may benefit from a similar analysis. Cultural signals in machines may relate to human codifications of average intelligence abilities for machines, whether it be fuelled by myths commonly at stake in science-fiction, or by trends currently related by the media [Haynes, 2003]. For example, our study of reinforcement learning for sonic exploration (Chapter 4) showed that humans perceived the agent as "light-headed", which accounts for cultural expectations on the expressive abilities of a machine.

7.3.2 Human-Machine Co-Expression

The blending of human expression with machine expression can create patterns of musical interaction that we propose to name *co-expression*. Of course, co-expression already happen during human performance with standard music instruments and interactive music systems [Jorda, 2005]. In line with our position, we believe that the notion of performance conveys too much of a bias toward expert performers from computer music communities. Co-expression, on the other hand, intends to provide a more general description of interaction between machines and (possibly non-musician) humans, in a way similar to human-machine co-adaptation [Mackay, 1990].

In this section, we discuss how co-expression may occur in human-machine music practice. We first (1) analyse interactive workflows designed in this thesis under the light of co-expression. We then (2) analyse how human-centred machine learning may specifically enable to learn machine expression from human expression. We argue that the framework of co-expression may help design more inclusive interactive music systems in terms of musical values and expertise.

Co-Expression as Emergent Cooperation, Competition, or Cohesion

In the course of this thesis, three interactive workflows were created: design through performance (Chapter 3), co-exploration (Chapters 4 & 5), and collective musical interaction mediated by an active learning machine (Chapter 6). We argue that each may be respectively analysed as emerging forms of co-expression: cooperation, competition, and cohesion. In the following paragraphs, we describe these emerging forms as strong concepts that could guide the design of human-machine co-expression *with* machine learning [Höök and Löwgren, 2012].

Cooperation occurs when two or more entities act together for a common objective. In music, cooperative acts between musicians may occur during composition—*e.g.,* discussion of musical ideas—and performance—*e.g.,* embodied communication of musical intention. In co-expression, cooperation may occur when the human perceives machine expression as jointly aligned with his or her musical expression. For example, our study of deep reinforcement learning for synthesis exploration (Chapter 5) related the emergence of co-exploration workflows that could be analysed as cooperative, where the machine takes actions on some VST with the goal of maximizing human feedback. Cooperation may be characterized in humans by the emergence of pleasant affective states due to the alignment of human and machine expression.

Competition arises when two or more entities strive for an objective that may not be shared. In music, competition may be analysed from the viewpoint of tension and resolution, which claims that emotion in music may be provoked by embodied evocations of divergent sonic states eventually coming to an agreement [Steinbeis and Koelsch, 2007]. In co-expression, competition may occur when the human attributes a divergent intention to the machine, that is, when local aspects of machine expression does not align with the global human expression. For example, our study of unsupervised learning for motion-sound mapping (Chapter 3) related the emergence of design through performance workflows that could be analysed as competitive, where performers strive to move differently to get specific sounds from the machine. Competition may be characterized in humans by affective states combining short-term unpleasantness with sustained arousal due to the challenging aspects of machine expression.

Cohesion occurs when bonds link entities of a group to one another and together. In music, cohesion may exist between diverse members of a specific musical community that share common cultural codes and a vision on music. In co-expression, cohesion may occur when the machine gently entrains the human to express in a certain way, without necessarily expressing through the same embodied quality, nor having a specific aim. It could be argued that cooperation and competition may be thought of as specific forms of cohesion between the human and the machine. Yet, cohesion

may be considered as a particular form of co-expression, due to its open-endedness in comparison with the two latters. For example, our study of active learning for collective musical interaction (Chapter 6) shows how the machine and humans form a cohesive group, which communicate through sound and visuals, yet do not share the same representational space nor having a specific musical aim. Cohesion may be characterized in humans by the emergence of positive emotion and mood due to the form of homeostasis [Leman, 2016] maintained by human and machine co-expression.

Learning Machine Expression From Human Expression

Machine learning techniques were central to the emergence of such co-expression workflows observed in this thesis. In the following paragraphs, we discuss how machine learning may actually learn to modulate human expression by creating models prone to adaptation, surprise, and incentive (see Fig. 7.3). We build on Leman's notions of human agency, prediction, and reward to describe human action data in terms of expression and support our claim [Leman, 2016].

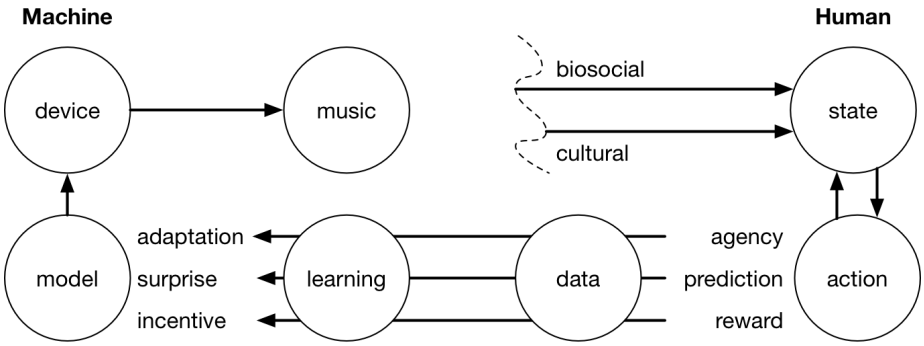


Figure 7.3: The strong concept of co-expression enables to qualitatively describe how learning from human data may enable machines to generate expressive musical behaviors based on human expression.

Agency is the capacity of an entity to act in a given environment. Music has been shown a sense of agency to humans in relation to their movement, for example when musical tempo aligns with human locomotion [Leman, 2016]. Machine learning have potential to augment human agency by creating a model that adapts in real-time to human data. For example, our design of unsupervised learning for motion-sound mapping (Chapter 3) created such space for agency by automatically adapting to human motion data to control sound synthesis. More generally, adaptation remains one of the main feature of machine learning techniques, which are designed to learn operational rules from example data [Bishop, 2006].

Prediction refers to the ability of a cognitive system to forecast a future event. It is often described as a key element in music interaction, where humans apply their own predictive sensorimotor model to the temporal unfolding of a given song to, for example, anticipate dancing movements [Leman, 2016]. Machine learning can potentially learn to manipulate such human predictive coding by purposely introducing musical

events that surprise human expectations. For example, our design of deep reinforcement learning for synthesis exploration (Chapter 5) enabled such predictive coding manipulation by taking exploratory actions on the basis of a human preference model learned from feedback data. Other algorithmic methods for surprise may be designed with machine learning, for example, by tweaking Bayesian inference to explicitly opt for unlikely decisions [Murphy, 2012].

Reward is the positive property that an entity ascribes to an event or an action. It is essential in the cognitive-motivational loop of humans and animals, notably to learn associations between events. Music listening, as well as music performance can be highly rewarding for humans [Leman, 2016]. Machine learning can potentially learn to manipulate rewarding associations between input and musical output to incentivize humans to rediscover the rewarding associations. For example, our design of active learning for collective musical interaction (Chapter 6) constantly manipulated the human reward model—corresponding to the association of a musical state with an audio effect—on the basis of collective performance data, so as to incentivize humans to play together. Incentive may be designed with other machine learning techniques that draw from rewarding mechanisms to learn models—*e.g.*, reinforcement learning [Sutton and Barto, 2011].

7.3.3 On Human Expression, Music, and Artificial Creativity

The notions of *machine expression* and *co-expression* enables to understand interaction in music systems from the perspective of *human expression*. As reported in Section 2.1.1, human expression, especially through embodied interaction with music, is central to the social and emotional development of human beings [Altenmüller et al., 2013]. Thus, supporting personal expression in interactive music systems may enable the development of creativity in humans [Frich et al., 2018]. Creativity is a general skill that may operate beyond artistic practice to a variety of activities, including scientific inquiry and technology design [Hart et al., 2017]. It may be supposed that the dynamic nature of co-expression, reinforced by the adaptive abilities of machine learning, may support *creative processes* characteristic of human music practices.

Many researchers, musicians, scientists, and designers have emphasized the fact that music is a *practice*, not a thing [Small, 1998]. As described in Section 2.1.1, music practices are necessarily of a dynamic nature, due to their situation in human, cultural, and societal contexts. To let new music practices thrive implies fostering the dynamic and improvisational nature of the creative process in humans [Canonne, 2012]. From a human-centred perspective, this means considering all possible dimensions of human interaction with music systems, encompassing technical as well as expressive elements—such as accounted with our strong concepts for co-expression. From a machine-centred perspective, this means designing technology that provide *space for expression* of possible solutions, rather than providing direct solutions—such as depicted by "metadesign" approaches to human-machine systems [Giaccardi and Fischer, 2008]. It is arguable that our human-centred approach to machine learning provided such space for expression, and as such, managed to take root in existing music practices.

In this sense, we prefer to use the notion of machine expression instead of that of *artificial creativity* when discussing machine learning applied to music practice. As described in Section 2.2.3, artificial creativity focuses on the building of autonomous machines capable of human-level creativity. We believe that such artificially-creative machines may drive the yearning to express one's creativity away from people. First, the autonomous abilities of these machines reduce human expression to a few parameters, thus hindering the emergence of co-expressive patterns of interaction. Second, the static generative behaviour of these machines go against the dynamic nature of creative processes at the core of existing music practices. We do believe that many researchers and developers in machine learning applied to music actually aim at designing musically-expressive, rather than artificially-creative, machines [Assayag, 2014, Esling et al., 2019]. Yet, we argue that their talking about "artificial creativity" creates an ambiguity regarding the scientific and musical values that they defend, especially toward fully-automated machines, typical of the latter approach.

Take the example of an artificial creativity machine able to automatically generate original orchestral symphonies by imitating the writing style of some famous composer. The biosocial signals expressed by such a machine arguably convey an embodied sense of superhuman performance, which relates to the instantaneous writing of a complete symphony. This may have expert musicians feel dispossessed of their hard-earned creativity [Surges and Dubnov, 2013], while also hindering non-musicians engaging in embodied interaction with music. On the other hand, the cultural signals expressed by such a machine arguably foster the notion of "artificial intelligence" and its related drawbacks, as discussed in Section 7.1.2. Whether resulting from personal or industrial motivations related to machine learning, or from a decent lack of contextualising, this happens to the detriment of alternative music communities—whose practitioners are disrupted by the introduction of such powerful technology—and for the benefit of mainstream music culture—whose members use, possess, and develop such artificially-creative machines. These points will be developed in the Post-Scriptum.

We believe that situated acts of creation can subvert such automated systems and techniques in favour of human expression and music practices. For example, human-centred approaches to machine learning enable to design expressive machines that balance human-level automation with human-level control, while accounting for the situated nature of the concerned music practices. Music practice supports situated appropriation of machine learning by members of given music communities, thus encouraging the development of other notions than artificial creativity to discuss machine learning [Assayag, 2014]. Artistic approaches to machine learning enable to nourish imagination on music to be made by such artificially-creative machines, especially by distorting their intended engineering [Williams, 2016]. Research, design, and applications led in this thesis clearly situates among the latter scientific and musical values.

8

Conclusion

8.1 Summary and Contributions

This thesis aimed at defending the approach of designing *with* machine learning for interactive music systems. To do so, it applied human-centred approaches to the design of machine learning, in contrast with engineering sciences approaches traditionally applied to machine learning. It investigated four musical tasks by situating them in their respective music practice contexts. It leveraged model prototyping to generate new interactive machine learning workflows with the concerned human users before engineering a final model. It applied four human-centred methods to research and design the musical interactions that took place between learning machines and humans. It adopted a reflexive approach to analyse the work done in the thesis and argue for the necessity to situate machine learning in human contexts to foster existing music practices.

Model Prototypes

We designed four model prototypes respectively corresponding to the four musical tasks investigated. Each of them implemented a machine learning technique in a human interaction setup, allowing to test novel interactive learning workflows with different types of data. Specifically, we applied the *Online Gaussian Mixture Model* to the musical task of motion-sound mapping by implementing unsupervised learning from human motion observations. We applied the *Sarsa* model to the musical task of sonic exploration by implementing reinforcement learning from human feedback and machine exploration. We applied the *Deep TAMER* model to the musical task of synthesis exploration by implementing deep reinforcement learning from human feedback, control, and intrinsically-motivated machine exploration. We applied the *Bayesian Inforation Gain* model to the musical task of collective musical interaction by implementing active learning from individual human performances.

Empirical Findings

We harvested empirical findings corresponding to the evaluation of human interaction with the respective four prototype models. The nature of these empirical findings related to that of the human-centred method employed to study interaction. Specifically, we found that performers projected themselves into imaginary situations interacting with unsupervised learning for motion-sound mapping. We found that humans were able to perceive differences in agent's behaviour with reinforcement learning for sonic exploration. We found that sound designers switched the lead with deep reinforcement learning agents for synthesis exploration under specific types of partnerships. We found that non-musicians reacted positively to the expressive behaviour enabled by active learning in collective musical interaction.

Music Dispositifs

We applied our four prototype models to the creation of several music dispositifs that we displayed in various public settings. The created dispositifs enabled to inquire machine learning applications to music in another type of situation than the lab-based studies of human-centred methods. Specifically, the *somasticks* are augmented drumsticks that use unsupervised learning to emphasize somatic expression in drumming practice. *Riding the Co-Explorers* is a single-level game that lets a single player guide the Co-Explorer in one single sound space. *Behavioral Matter* is an interdisciplinary workshop that applied the Co-Explorer to the design of robotic objects. *ægo* is an improvisational computer music piece with interactive sound and image for one performer and the Co-Explorer. *entrain* is a public installation that lets humans express collectively in music in collaboration with an active learning loudspeaker.

Theoretical Perspectives

We reported on theoretical perspectives on machine learning applied to music by adopting a reflexive approach on our work. These perspectives enable to rethink machine learning applied to music in terms of human goals and values. Specifically, the artifacts called *model prototypes* help envision possible designs of machine learning in situation with human users before fully engineering them. notion of *machine expression* supports the general description of human perception of machines in music practice, regardless of humans' expertise in music. The notion of *co-expression* enables to conceptualize the emergence of musical interaction between (possibly non-musician) humans and (possibly learning) machines. Overall, the notions provide actionable alternatives to artificial creativity approaches to the engineering of machine learning, which currently drive most of machine learning applications in music.

8.2 Directions for Future Work

While our interdisciplinary approach enabled original contributions, it also left several research questions open. We conclude the present dissertation by outlining

human-centred, machine-centred, and musical perspectives for future work to be done.

8.2.1 Human-Centred Perspectives

Designing with machine learning for interactive music systems had us explore four situated music tasks instead of focusing on one sole music practice. This choice let many questions open related to the studied human users, be they musicians or non-musician people.

On Musicians

Two of our studies—Chapters 3 and 5—situated the research and design of machine learning in specific music practices—motion-sound mapping and synthesis exploration, respectively. Yet, more time could have been spent working closely with the corresponding *expert musicians* to get a better understanding of their practices—performance and sound design, respectively. This understanding could notably enable to better evaluate the creative process of musicians in the appropriation of machine learning [Fiebrink et al., 2011]. Concretely, longitudinal studies could be led with sound designers, individually using our deep reinforcement learning model prototype over longer periods of time, to observe interaction with machine learning more precisely. Also, field studies could be led with performers, observing them interacting with existing interactive music systems in their own creative environment, to get an understanding of their musical practice that go beyond the designed machine learning artifacts.

On Non-Musician People

The two remaining studies—Chapters 4 and 6—situated the research and design of machine learning in more general music practices—sonic exploration and collective musical interaction, respectively. Yet, more time could have been spent studying how *non-musician people* can develop their creativity through music practice. These studies could in turn enable the design of machine learning that support education and inclusion purposes [Thompson et al., 2019]. One of the paths suggested by this dissertation would be to study interaction between humans and machines from the perspective of co-expression. Longitudinal studies could be led, investigating processes of expressive interaction between non-musicians and machines, to observe the development and crystallizing of strategies in relation to some musical task—such as sonic exploration. Alternatively, qualitative studies could be led in public settings to better understand how non-musicians perceive machine expression in open-ended musical activities—such as collective musical interaction.

8.2.2 Machine-Centred Perspectives

Designing with machine learning for interactive music systems had us explore four machine learning techniques instead of focusing on one sole machine learning model.

This choice let many questions open related to the studied techniques, notably regarding the engineering of models, and the quantification of machine expression.

From Model Prototyping To Model Engineering

Our four studies built on model prototyping to ideate applications of machine learning with human users. Only one—Chapter 5—benefited from model engineering to fine-tune the design of machine learning for the corresponding musical task—namely, synthesis exploration. Yet, further iterations in the design processes could have been done to go beyond one-step prototyping to reach *model engineering* in all four studies. Such reachings may enable the creation of design patterns that explicitly consider the specificities of machine learning in human interaction contexts [Beaudouin-Lafon and Mackay, 2009]. For example, several model prototypes could be explored in parallel to assess their suitability to human interaction. Larger data bases of human interaction could be built to support the pre-engineering of models through offline learning. Human-centred optimization criteria may be developed to engineer the final models in relation to some situated musical task—such as the intrinsically-motivated exploration criteria in Chapter 5.

Designing Machine Expression

Reflexive approach on our four studies—in Chapter 7—enabled to frame machine expression as a qualitative notion that describes musical interaction between (possibly non-musician) humans and (possibly learning) machines. Future work could explore a reverse approach consisting in *designing machine expression* in interactive music systems. Such an approach may combine qualitative and quantitative descriptions of expression to design machines that go beyond the false divide between the digital and the physical [Bianchini and Verhagen, 2016]. Other types of data could be investigated to support expression in interactive music systems. For example, physiological data could be studied to sense, or even measure, musical expression in humans. Complex environmental data could be leveraged to support artificial implementation of expressive features in machines. Machine learning could be used to explicitly model expression in humans and machines, as well as to support algorithmic description of co-expression.

8.2.3 Musical Perspectives

Designing with machine learning for interactive music systems had us apply four human-centred methods and lead five applications in real-world musical situations instead of focusing on one sole methodology. This choice opens many perspectives related to the methods to be used for, and the nature of, music research.

On Human-Centred Methods

Our four studies used human-centred methods to build scientific knowledge on musical interaction. These four methods allowed us to somehow synthesize perspectives

from the many disciplines involved in the work, such as Human-Computer Interaction, Machine Learning, or Music. Yet, further investigations could be led to better understand how these *human-centred methods* actually influenced the design of machine learning for the situated musical tasks. These understandings could enable to frame new methodologies for the application of Human-Computer Interaction methods to the field of Music [Wanderley and Mackay, 2019]. For example, the design-oriented method used in Chapter 3 created an "idiosyncratic" formalization of unsupervised learning for motion-sound mapping, which could have been framed differently by other designers and performers of interactive music systems. Alternatively, the experimental method used in Chapter 4 directed the formalization of reinforcement learning toward a "cartesian" form to enable the controlling of the many environmental factors of the experiment.

From Applied Research to Practice-Based Research

Our five applications in real-world musical situations let us share the experience and the conceptual ideas of our interactive music systems with people and society in general. These creative applications somehow complemented the scientific findings leveraged through our human-centred studies. On another note, more effort could be put to go beyond such applied approaches to embrace *practice-based research* in music. Such an approach could enable to situate inquiries of machine learning from other research communities delineated by real-world music practices, as it has recently been done by composers of contemporary music [Ghisi, 2017] and experimental music [Williams, 2016]. For example, our improvisational piece *ægo*—presented at the end of Chapter 5—could benefit from longer processes of research and creation to eventually exist as actual artwork outside the scientific research community. We believe that such situated, practical, and sensible approaches to music research should never be disregarded when studying, designing, or evaluating interactive music systems.

Post-Scriptum: Music Research Through Dispositif Design

In this chapter, I present my personal reflections on the research approach adopted in the thesis. In accordance with my advisor, I chose a post-scriptum format, which allows me to discuss ideas in a less formal setting, using references and a first-person narrative to support my personal statement rather than a general analysis. Rather than discrediting possible research approaches, my wish is to make my perception of the field audible as doctoral student, with the hope to discuss a common ground for music research to be done in a near future. I hope to hear back from other computer music researchers, practitioners, and students about it.

Doctoral Research in Science Applied to Music

The present dissertation was written in the frame of a doctoral research in Computer Science applied to music. As in most formal science disciplines, Computer Science doctoral schools create pressure for ending the doctorate in three years. This relatively short period of time creates a situation where doctoral students are pushed toward being productive in the activity of research. The productive impetus materializes into an urge to rapidly *publish* in academic journals and conferences. Even if not explicitly stated in doctoral schools decrees[1], publications are an important modality of evaluation for theses in the academic community. For example, it is a common practice to seize the quality of a thesis by having a look at the list of published papers. Alternatively, the impact factor gained from publications remains one of the main criteria for the pursuit of an academic career, be it at a national or international scale.

Publishing in academic journals and conferences almost always implies adopting scientific *methods*. Methods create a common framework to discuss empirical observations between peers, and have thus characterized the development of science for centuries. Recently, human-centred methods were object to critiques [Vandenberghe and Slegers, 2016]. The key point is that simply involving humans in the loop does not

[1] https://www.legifrance.gouv.fr/affichTexte.do?cidTexte=JORFTEXT000032587086&categorieLien=id

account for the resulting design to be the right for the studied "users". In our case, the relying on human-centred methods encouraged the *application* of computer science to music. That is, they supported the application of scientific processes to solve practical problems in music. For example, in Chapter 5, we used a user-centered design method to adapt the design of reinforcement learning to the musical task of synthesis exploration. Yet, we did not discussed our initial choice of reinforcement learning with the studied expert sound designers. It is possible that a completely different technique would have emerged if we let sound designers discuss the initial design problem with us—if there were any real problem to solve.

Methods focus on creating *substantial* accounts of knowledge, relying on standardized languages to communicate results to peers. As stated, a first language is that of problem specification. Specifying a problem enables to frame a research question in relation to a research method, and to look for a solution. A second language is that of measurement. Quantification—most of the time materialized into mathematic computation—constitute the main modality of measurement to assess the efficiency of some solution in the frame of some method. A third language is that of wording. Most academic journals and conferences rely on English to communicate knowledge, as well as on templates that standardize the structure and format of knowledge. The dissertation format of the Computer Science thesis also highly constrain the way to relate musical research. To sum up, substantial accounts may optimize the production of knowledge in Computer Science, but in turn, may tend to undermine form as sensible component of musical knowledge.

I suggest that the "substance over form" paradigm may originate from the *disciplinary* educational establishment provided in France. Obtaining a scholarship to do doctoral research in Computer Science almost inevitably requires having a Master's Degree in a formal science discipline—*e.g.,* mathematics, engineering, computer science, or physics. I argue that such trainings push doctoral students to be productive in their research within the *norms* of the scientific approach. Reciprocally, the academic status of music in France reinforces the disciplinary framework for research. For example, the recently-created doctorate in Music[2] is intended to "high-level composers"—*e.g.,* already having professional experience in music—, and focus on one formal music practice—*i.e.,* composition. I claim that this academic context may not encourage scientifically-trained doctoral students to go beyond applied research to explore practical and sensible approaches to music research. In the following paragraphs, I explain why I believe that research through design approaches may support the negotiation of these scientific and musical norms in doctoral research.

Music Research Through Design

Design is a creative activity that focuses on the crafting of systems adapted to humans. Design as a conceptive activity field borrows techniques from many scientific disciplines such as engineering sciences, computer science, psychology, or biology. Yet, as Nigel Cross put it, design should be not be considered as a science, but rather

[2] https://www.ircam.fr/transmission/formations-superieures/doctorat-de-musique/

as a discipline in itself, characterized by "designerly ways of knowing" [Cross, 2001b]. This specific way of knowing of designers may be clarified by Alain Findeli's notion of *project* [Findeli, 2004]. While scientists tend to consider the world as an object that they study—dissociating them from the world as subjects—, designers tend to consider the world as a project—*i.e.,* they situate themselves as subjects in the process of studying an object. Design approaches may thus be suited to music, which is, as we described in Chapter 2, a socially- and culturally-situated process.

Practice generally prevails over method in design. While scientific methods aim at validating results and guaranteeing their reproduction, design practice does not, or should not, be repeatable, as a situated act of creation [Cross, 2001b]. As such, design practice is efficient to solve "wicked problems", which, as Horst Rittel coined it, are socially-rooted problems that are better suited to iterative reframing than to methodical reasoning [Rittel, 1972]. The practical aspect of design makes it close to music, which is also defined by its practical and action-oriented processes [Leman, 2008].

Form may be as important as substance in design practice. The created artifacts may rely on aesthetics, composition, or style to exert their intended functions. These subjective and creative dimensions of artifacts may add to their objective and scientific dimensions, such as their technical structure and organization [Archer, 1979]. Crucially, evaluation of the form component cannot be reduced to the measurement of a numerical variable in a controlled setup. Rather, it is the cultural or industrial applications that may testify of the success of a given design in a given real-world context. Design may thus be suited to the building of interactive music systems, which similarly rely on technical structures to create sonic forms that are to be interacted with by humans in real-world musical situations.

Design *research* may have as many materializations as there are of design communities [Vial, 2015]. For example, the field of design science aims at "scientizing" the knowledge of design, typically by building design methods as well as publication venues to standardize the design discipline in a way similar to science [Cross, 2001b]. Yet, many designers consider that design research must include elements of practice-based research, in a way similar to artistic research [Candy, 2006]. For example, speculative design uses design artifacts to anticipate technological deviations and harness public reactions toward it [Auger, 2013]. Critical design use design as an artistic medium to research the process of design by inquiring designers on their own practice [Bardzell and Bardzell, 2013]. In parallel, Francophone communities have essentially focused on industrial and aesthetic applications of design, arguably because of a lack of an appropriate context for academic research [Monjou, 2014]. To borrow from Annie Gentès [Gentes, 2017], this *indisciplinarity* of design may be another trait of similarity with music, which is also praised for its indisciplinary in Francophone computer music communities [Sèdes et al., 2019].

I argue that carrying music research *through* design may provide opportunities to reconsider both scientific and musical norms in interactive music systems. Research through design is a practice-based approach to design research, consisting in (1) making projects in real-world situations, and (2) relate on the framing of the problem itself—*e.g.,* the balance that was made between the intersecting and possibly conflict-

ing perspectives [Zimmerman et al., 2007]. I argue that the practice-based contributions enabled by research through design approaches may complement the applicative contributions of human-centred approaches in the research on interactive music systems. I believe that the reflective practice of translating tacit design knowledge into words may globally benefit music research [Schön, 2017]. In the next section, I introduce the concept of *interactive music dispositif* to frame such *music research through design* approach.

From Interactive Music Systems to *Interactive Music Dispositifs*

The term "dispositif", originally introduced by French philosopher Michel Foucault, names the set of normative elements that may condition a given human activity [Agamben, 2009]. Dispositifs include both *endogenous* and *exogenous* aspects of interactive music systems. That is, they encompass technical knowledge on their internal implementation, along with the imaginary aspects that they may be transmitted to their external environment. While "system" could also be used for a similar definition [Burnham, 1968], I claim that the term conveys too much of a technical meaning in music communities (see Chapter 2 for an overview). Although popular in the fields of art and design [Bianchini and Verhagen, 2016], dispositifs remain under-exploited in the field of music, with a few recent exceptions in the field of sound art [Sinclair, 2018].

I introduce the notion of *interactive music dispositif* to explicitly consider the scientific and musical norms that may condition human knowledge, actions, and feelings toward interactive systems. I argue that *music research* can be led through *dispositif design*, that is, through the technical realization of aesthetically-functioning artifacts that challenge these cultural norms. In the next two sections, I provide a critical assessment on the scientific and musical norms that may determine interactive music dispositifs, focusing on machine learning as a scientific discipline. Rather than offering a generic baseline for dispositif design, my wish is to share personal thoughts based on the practical experience I gained through real-world applications of the dispositifs created in my doctoral research. I hope that this initiative may engage a more global reflection with other music researchers.

Norms in Music

As a culturally-learned practice, music may encapsulate several norms. In the next sections, I propose a description of these normative elements, which I grouped in terms of musical (1) knowledge, (2) action, and (3) feeling (see Figure A).

Musical Knowledge

Language conveys norms related to the description of music. Notation may perpetuate cultural knowledge on structures and grammars of music, potentially restraining sound-oriented musicians and non-musicians from interacting with a music dispositif

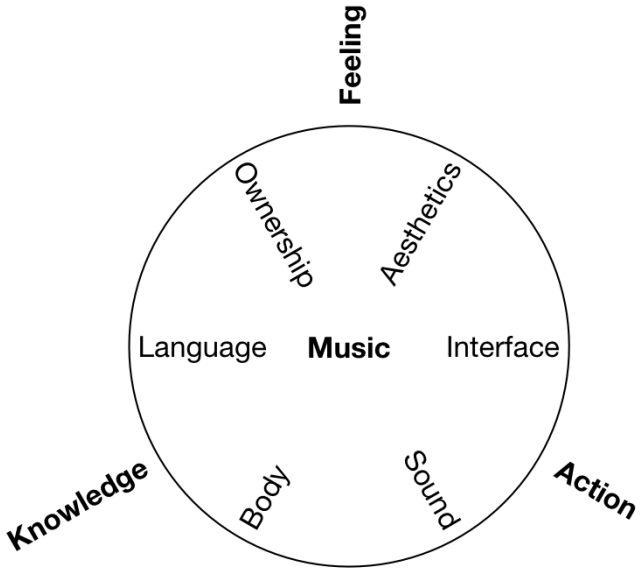


Figure A: Normative elements related to music.

[Magnusson, 2019]. Also, technical terms, such as parameters' names of synthesizers, may act as a disincentive for non-musicians to appropriate music systems, as our conceptual map showed (Chapter 4).

The human *body* conveys norms on the training required to practice music. Physical techniques related to music practice are known to require considerable time to be learned and mastered by humans [Palmer, 1997]. Skeuomorphism tend to perpetuate nonlinguistic elements related to physical techniques in the digital domain [Lindh, 2018]. The *somasticks* (Chapter 3) are an example of dispositif that conveyed assumptions on the bodily techniques required to play with it.

Musical Actions

Sound conveys norms on the actions to be taken by performers and listeners of a dispositif. First, the range of timbres of an interactive music dispositif may condition human actions in music making. For example, the *entrain* dispositif (Chapter 6) only enabled humans to play a restrained set of prerecorded drum samples. Second, the expressive qualities of sound may condition human actions in music listening [Leman, 2016]. For example, the rhythmic loops generated by the *entrain* dispositif may leverage cultural norms on the way one should act when listening to it.

Interfaces of dispositifs reinforce norms on the musical forms to be created with sound. First, the appearance of interfaces may orient human actions in music making, as made famous by the notion of affordance [Gibson, 1977]. Second, the technical implementation of dispositifs may restrain human control over the processing of sound in time. For example, interactive music dispositifs dedicated to music production, such as Ableton Live, highly encourage humans to create aesthetically-confined

music—*e.g.,* the MIDI-based implementation combined with the clip functionality. On the other hand, the *Riding the Co-Explorers* dispositif (Section 5.8) managed to appropriate the norms in music interfaces through the use of video game.

Musical Feelings

Aesthetics tend to standardize human feelings related to the perception of music. For example, human interaction with a synthesizer may be implicitly influenced by aesthetic elements assessing the good and bad usages of some parameters. Music institutions have an active role in defining the aesthetic elements of interactive music dispositifs. Depending on the will of institutions, some aesthetics or music genres may be criticized to base a dominant norm in relation to one given music practice [Born, 1995]. Or, they may be valued to sustain change and favour the emergence of diverse music practices [Dryhurst, 2019].

Ownership tends to determine human feelings on the right to make music. In the frame of the thesis, we made an effort to spread our interactive music dispositifs for free using many different public dispositifs—*e.g.,* educational workshops, open source projects, or public exhibitions. In a more general frame, individual ownership of techniques related to music may create elite art forms, partly defined through socioeconomic privilege to use technology [Hyde, 2010]. Alternatively, industrial ownership may displace the value of music from that of a cultural to a monetary object [Wikström, 2013].

Norms in Machine Learning

In the frame of my doctoral work, I now describe the norms encapsulated by machine learning in relation to Computer Science. I attempt to describe these normative elements by grouping them in terms of human (1) knowledge, (2) action, and (3) feeling (see Figure B).

Human Knowledge

Data conveys norms on the way humans may create knowledge of the world. Measuring data implies possessing a dispositif capable of detecting a phenomenon in the world with sufficient accuracy. Then, collecting and reporting data implies one or several humans leading observations in the world to measure phenomena. Thus, data is far from being a neutral element of knowledge, but rather a situated, partial, and constitutive one. Researchers in the Humanities proposed to rename it *capta* (from the Latin capere, "to take") to emphasize the fact that data has been taken from a phenomenon by an observer using some measurement dispositif [Drucker, 2011].

Data sets may reinforce the norms encoded in data at larger scales of knowledge. For example, data sets built from popular music [Bertin-Mahieux et al., 2011] may reinforce the power relationships between musical genres, notably by enabling large-scale comparisons only based on a few musical features. On the other hand, human-centred machine learning approaches leverage small-scale data sets [Gillies et al., 2016]. These approaches enabled us to switch data sets from normative to evolutive

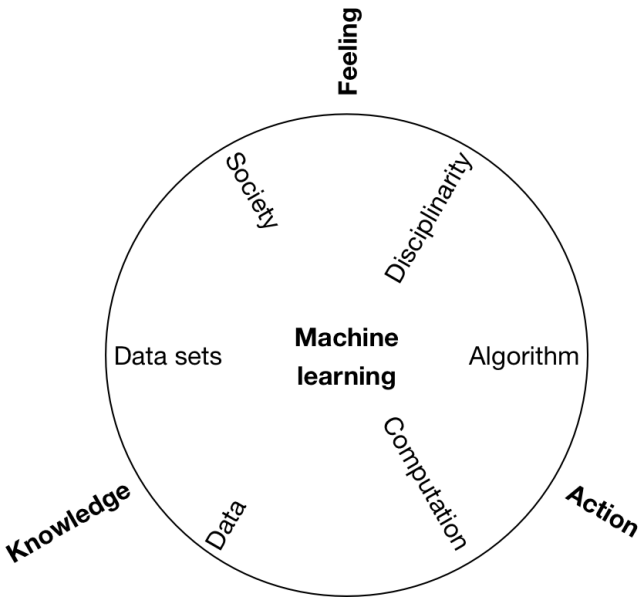


Figure B: Normative elements related to machine learning.

technical components [Giaccardi and Fischer, 2008]. This evolutivity was witnessed in the *Behavioral Matter* workshop (Section 5.9), where small data sets were shown to support robotic object design as well as human synthesis exploration.

Human Actions

Computation encapsulates norms on the machine learning actions of creating, and reasoning about, data knowledge. First, many machine learning computations rely on inferential statistics, which produce probable, yet inevitably uncertain new knowledge [Mittelstadt et al., 2016]. Second, computation may enable correlations between two variables, but may be insufficient to prove a causal relationship [Illari et al., 2011]. Explainability of machine learning is still subject for research to address the increasing complexity of computations—*e.g.,* in the many layers of deep neural networks [Richardson et al., 2019].

Algorithms create new norms that condition human actions at the scale of society. First, machine learning algorithms enable automatic definition of decision-making rules. For example, a classifier may learn to order data in terms of categories to which they belong, but may also produce new categories based on its mathematical construct [Domingos, 2012]. Second, algorithms can have a transformative effect on the way humans conceptualize the world [Mittelstadt et al., 2016]. The most basic algorithmic processes—*e.g.,* gathering data and using algorithms to classify entities—thus determine new human actions that may not have been thought of without such dispositifs. The improvisational performance *ægo* (Section 5.10) aimed at illustrating such directing of human actions by machine learning algorithms.

Human Feelings

Disciplinarity of public scientific institutions tends to standardize the way machine learning is taught and researched on. The last decade has witnessed a wave of strategic and political plans related to "Artificial Intelligence" all over the world [Cath et al., 2018, Villani et al., 2018]. Public universities, research organizations, institutes, and laboratories are allocated funding grants to research and develop "Artificial Intelligence". This governmental impetus may create pressures on academics from the engineering sciences as well as the humanities to apply machine learning to their work, letting few space for critical reflection. The current thesis work, I believe, have remained critical toward this normative wave, by adopting an interdisciplinary approach to the design of machine learning.

Society tends to determine human feelings on machine learning and computer science in general. "Artificial Intelligence", as the last intensification of the digitization phenomenon, offers tremendous opportunities to market goods and services to the world [Crawford and Joler, 2018]. Machine learning applications thus grapple with capitalism, and the globalized and privatized system that goes with it [Deleuze, 1995]. These elements may influence how people think about machine learning, as we witnessed in Chapter 7. Along with other academics [Caramiaux et al., 2019], I believe that the upcoming prevailing of private industries in creative fields should be addressed by public research communities, which include music research communities.

All norms stated above reflect my critical assessment of endogenous and exogenous aspects of interactive music systems. I invite other researchers to reflect on their own practice to contribute to a more general definition of interactive music dispositifs.

Example Projects

In this section, I show how dispositif design enables to challenge each of these normative elements. I rely on three example projects of practice with music and machine learning, led in parallel of my doctoral research, to illustrate this claim:

- *Sound Control* was an action research project started in 2016, led by Rebecca Fiebrink, in which I participated as pre-doctoral researcher at Goldsmiths University of London [Scurto, 2016].
- *Paroles d'Exil* was a research and creation project ended in 2017, led by the composer Mélanie Egger at GMEM-CNCM / Friche la Belle de Mai, on which I worked two years as computer music designer.
- *{Lutheries}*^{2} was an outreach project carried in 2018, initiated by Lutherie Urbaine, that I directed during a one-year volunteering as education assistant on science and music.

Figure I schematizes how the three projects respectively injected human values in the above-mentioned scientific and musical norms.

Sound Control

The *Sound Control* project was initiated by Simon Steptoe, together with Louise Tyrrell, who runs the Musical Inclusion Programme for the Northamptonshire Music and Performing Arts Trust (NMPAT). They notably collaborated with the researcher Rebecca Fiebrink, along with music educators and music therapists associated with a community music centre's Musical Inclusion programme. Programme members were interested in more flexibly customise digital instruments for disadvantaged children they worked with—including but not limited to children with physical and learning disabilities. All had experience supporting youths' acoustic music-making (e.g., singing, using simple percussion instruments) and using simple switch-based controllers to trigger pre-recorded samples. They were interested in how they might use bespoke sensor-based instruments with youth in their programmes.



Figure C: Picture of an inclusive ensemble[3] performing with *Sound Control*.

Our team led eight workshops (approximately 1-3 hours each) with Musical Inclusion programme personnel and other music therapists and educators from the local community (see Fig. C). Early workshops showed participants demonstrations and videos of existing approaches to creating bespoke musical instruments (e.g., the British Paraorchestra, Wekinator[4]), then engaged participants in brainstorming activities. In later workshops, we taught participants to use prototype technologies developed for the project, such as Grab-and-play [Scurto et al., 2016], then elicited feedback about them. In parallel, our team attended seven classroom and workshop sessions in which two practitioners from the programme worked with children with disabilities in a school (five sessions) or community centre (two sessions). A typical session involved one teacher or therapist working one-on-one with 3-6 children for 10-30 minutes each, with members of our team helping with technology as needed. Additionally, the two practitioners began using the Sound Control software[5] independently after the fourth of these sessions, and they sent us bug reports, feature requests, and updates on their usage by email.

Sound Control enabled (1) sonic **expression** of non-musician people through machine learning computations, and (2) **accessibility** of embodied music-making using sensor data (see Fig. I, bottom). First, the development of the Sound Control software

[3]https://network.youthmusic.org.uk/nmpat-our-digital-ensemble-pilot-session
[4]http://www.wekinator.org/
[5]http://soundcontrolsoftware.com/

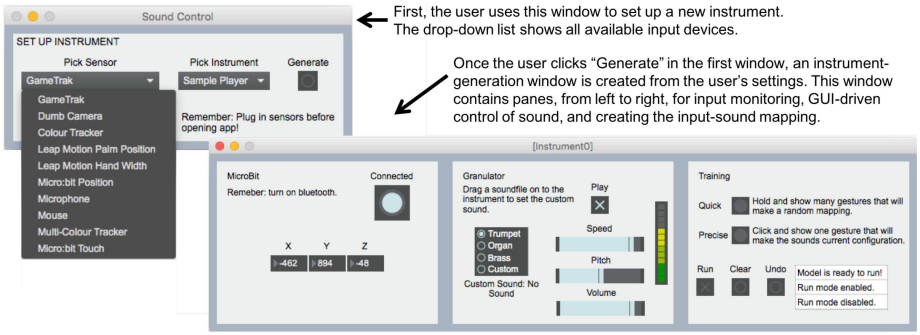


Figure D: User interface for the *Sound Control* dispositif (taken from [Thompson et al., 2019]).

let us learn new insight on the requirements needed for musical expression with machine learning. Human criteria—such as recognizing and exercising agency, encouraging moving and listening, and supporting social aims—were elicited by the design team, which pushed us to opt for simple sonic interactions realized by straightforward machine learning computations. Rather than building adapted body-based interfaces, the resulting music dispositif helped carers create diachronic forms, that is, short-lived data configurations reflecting the current state of interaction at an experiential level [Bourganel et al., 2016]. Second, the process of collaboration with music therapists, teachers, and children improved accessibility of music and machine learning. Using interactive supervised learning enables to create new sensor- and vision-based musical interfaces based on data examples of motion and sound, making interface building fast and accessible to people without programming or engineering expertise. The fact that Sound Control continues to support the therapists and teachers in their workshops, and enabled the children with disabilities to take part in public music performances, let us think that co-design may be a promising approach to frame machine learning for music.

Paroles d'Exil

Paroles d'Exil (2017) is a music installation co-created with Mélanie Egger during a residency at the *GMEM-CNCM / Friche la Belle de Mai*, Marseille, France, in collaboration with the *Sound Music Movement Interaction* group of *IRCAM* (see Figure E). The residency was led in the context of a Master program in Music from *Aix-Marseille Université*. Mélanie is a composer and pianist who took this two-year program to carry a research and creation project, named *Remembrance Fragments and Documentary Music*. Her project, supervised by Christine Esclapez, crossed historical facts, philosophical writings, musical and contemporary artworks to nurture a personal reflection on the human notion of time. Its realization took two forms: a Master's thesis on the one hand, and a musical artwork on the other hand, named *Paroles d'Exil*, for which she kindly asked me for a collaboration.

For this artwork, Mélanie primarily wished to work with archival voices of the Algerian War period—in relation with her personal history. Prior to collaborating, she



Figure E: Pictures taken during the restitution of *Paroles d'Exil*.

thus harvested a corpus of sounds from radiophonic emmissions with corresponding dates and location. Then, the highly-recursive process of collaboration included: discussion to become acquainted with Mélanie's musical concerns, demonstrations to show the technical possibilities offered by IRCAM's tools, interpretation of their scientific basis for Mélanie to grasp theoretical links between the tools' technical manipulation of sound material with her personal reflexions, creation of sonic excerpts to grasp experiential aspects in relation to the intended aesthetics of the piece, and final design of the music dispositifs during the residency. In close link with her memoir, we converged on a music installation format in three parts—*Histoire*, *Mémoire*, and *Oubli*—, each comprising a music dispositif treating archival voices in an appropriate manner. The second dispositif, *Mémoire*, leveraged machine learning as a central conceptual and technical element of its design.

Mémoire continually generates sequences of syllables through loudspeakers fragmented in location. The syllables, originally extracted from a corpus of archival voices dated from the Algerian War period, are reorganized as sequences in real-time by a machine learning algorithm, who creates new narratives through the recorded voices. Concatenative synthesis is used to decompose original sound recordings into separated syllables based on their spectral content (see Fig. F, left). A factor oracle is used to create a sequence model that relates suffix and affix probabilities of transition based on the spectral similarities of syllables originally presented in the corpus. Navigation in the sequence model is made autonomous by the oracle. Remembrance is imposed on us as a split, selective, hardly decipherable entity, in a constantly-evolving present.

Paroles d'Exil enabled (1) **appropriation** of machine learning algorithms to create a new interface between humans and sound, and (2) **reflection** on design as a creative discipline that contributes to musical aesthetics (see Figure I, upper right). First, the development of the *Mémoire* dispositif led to a creative use of machine learning. The algorithm that we used, namely CatOracle [Einbond et al., 2016], was originally designed for human-computer improvization. Recursive developments and discussion with Mélanie enabled us to transform this algorithm as a new music interface—*e.g.,* as an autonomous music installation—, appropriating its technical affordances to the music artwork. Second, the process of collaboration between Mélanie—the composer—and myself—the designer—let her reflect on the creative role that a designer may have in relation to music. After these two years of collaboration, she eventually considered

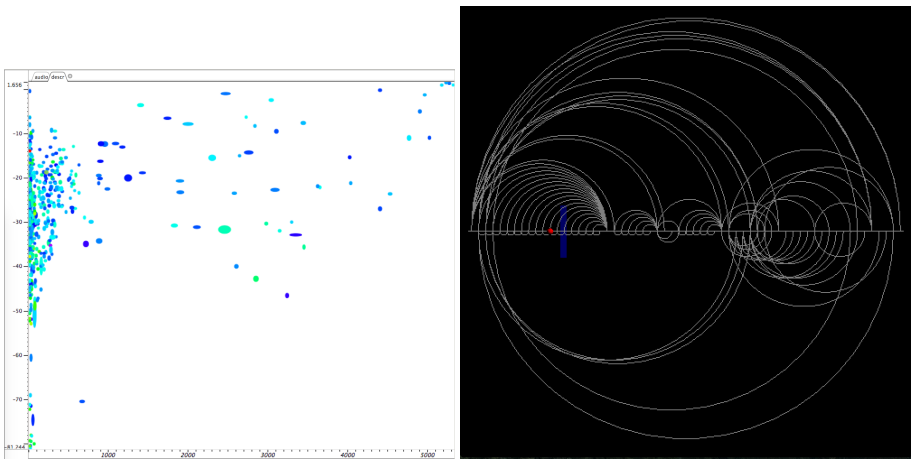


Figure F: Control interface for the *Mémoire* dispositif (built on CatOracle).

myself to be credited as artistic co-creator for the music installation, as the technical realization had a deep link with her theoretical and aesthetic reflections.

{Lutheries}²

{Lutheries}² is an exhibition conceived with *Lutherie Urbaine* and Laurent Le Gall (see Fig. G). It was intended to introduce people of all ages to the relationships between science and music through interaction with acoustic, electronic, and digital music dispositifs. The visit tour consisted in seven areas. A first area—*History*—introduced the links between technique, science, and technology by recounting the evolution of lutheries. The following areas recounted this historical course through the discussion of several notions linking science with music—*Harmony*, *Acoustics*, *Timbre*, *Electronics*, and *Digitalities*. The tour finally came out on an area dedicated to play—*Rhythm*. Illustrated, easy-to-follow panels were displayed along the interactive journey to elaborate on the technical functioning of the instruments. An educational kit accompanied the exhibition and described several fun activities conceived for each of the notions developed in the areas.

As diverse as the free-to-play music dispositifs of the exhibition were, their common feature was their being made of recycled objects. The *Lutherie Urbaine* association indeed accumulated twenty years of experience in the recovery of waste material to craft unique music dispositifs. Examples of waste included materials—*e.g.,* plastic tubes, wood boards, metallic barrels, glass bottles, or electronic circuits—but also everyday objects—*e.g.,* tennis rackets, beverage bottles, or electronic toys. These music dispositifs were often crafted in the context of musical artworks, during which they may stand as central to the narrative. Through these artworks, *Lutherie Urbaine* wished to sensitize the public to sustainable development. Many musicians and luthiers that pertained to the association led education outreach activities in public elemetary and secodary schools, sharing their unique crafts to a diversity of pupils, who could bring their own abandoned objects and bring them back to life in music.



Figure G: Pictures of the *{Lutheries}*^{2} exhibition at *Bagnolet, France*.

For the *Digitalities* area, we conceived a new dispositif, based on CataRT (see Fig. H). The dispositif presents the public with several recycled objects put on a table. Behind the table, a screen and a loudspeaker stand still, but are not accessible to use by the public. As soon as the public attempts to manipulate the recycled objects, the digital dispositif may react by generating sound along with a screen animation. Sounds are created on the basis on the frictions made by the object on the table. A piezoelectric sensor is put at the surface of the table, converting the diversity of vibrations made by material objects into digital data. This data signal is sent to a shallow machine learning algorithm—k-nearest neighbor—, which computes the similarity between the object's vibration and a data set of environmental sounds recorded in various locations—represented as points on the screen. The most similar sounds are then played by the dispositif, who accounts for the expressive cues encoded in the materiality of the objects, jointly with the physical manipulations made by the public.



Figure H: Picture of the music dispositif designed for the *Digitalities* area.

{Lutheries}[2] enabled (1) **diversity** of musical languages through an environmental data set, and (2) **sustainability** of machine learning through a public music dispositif (see Fig. I, upper left). First, the development of the dispositif in the *Digitalities* area led to the creation of a data set made with environmental sounds, rather than instrumantal sounds, or music songs. This creates a diversity of languages for music, linked with the idea of matter through their sonic affordances, and echoing approaches from thing-centered design [Giaccardi and Fischer, 2008] as alternative to human-centred approaches. Second, the outreach collaboration with *Lutherie Urbaine* encouraged sustainable views on machine learning and music. Rather than using deep architectures, our dispositif relied on a shallow machine learning algorithm. Not only was it cheaper in terms of computational cost, but it also promotes alternative uses of machine learning that may differ with those commonly thought of in society. Also, the context of public exhibition let people freely experience with the dispositif and the original sonic interactions that it fostered in relation to the recycled objects. Sustainable computers, speakers, and sensors may be used in next versions of the dispositif.

Music (or) Design Practice?

The projects related above open questions related to the *nature of the practice* on interactive music dispositifs. A first option would be to consider them as being part of a *music practice*. Yet, they were not led through the prism of composition, which is, as mentioned in the beginning of this post-scriptum, the main modality of music practice in the French academic music community. Sound art may be a more legitimate category, as a relatively recent artistic practice traditionally not associated with music [Wishart, 1996]. Yet, I believe that interactive music dispositifs should remain attached to music as a central field of work, for example by having other protagonists such as computer music designers, performers, authors, or teachers, to be recognized as equivalent to composers [Zattra and Donin, 2016].

On the other hand, recognizing the projects as *sound design practice* is incorrect to my mind. Sound design focuses on the use of sound to effectively fulfill some functions, representations, or signalling [Misdariis, 2018]. Yet, the three projects, as well as the notion of dispositif, put a stronger emphasis on human interaction to realize these functions than on sound itself. Sonic interaction design may be an alternative path to characterize the practice [Franinović and Serafin, 2013]. Yet, the emerging field lacks the musical dimensions of interaction, which directly deals with the cultural roots of music, and that I wished to emphasize with the notion of interactive music dispositif.

To borrow the expression of Norbert Schnell[6], I would like to call *music design* the act of conception and reception of interactive music dispositifs that supports musical communication. Music design values the social and cultural dimensions of music [Small, 1998, Frith, 1998, Herbert, 2012] over the practice of music composition. Beyond sound or embodied interaction, music design is concerned with *musical communication* supported by dispositifs [Miell et al., 2005, Leman, 2008]. Musical com-

[6]Norbert recently created a Master program in MusicDesign: https://www.hs-furtwangen.de/en/programmes/musicdesign-master/

munication is the process of expressive interaction that take place between entities that share an acoustic environment—be they humans, machines, animals, or things [Altenmüller et al., 2013, Leman, 2016]. In this way, the practice of music design may tend to shift the norms with which we listen and interact with our shared environments, in a current context of technofetishism [Latour, 1996]. The three above example projects and the corresponding dispositifs may support such music design practice, focusing on communication between humans and diverse elements of their environment—encompassing people with disabilities, cultural memories, materials and living things.

Overall, I believe that the practice of music design may in turn greatly benefit *music research*. Similar to other design fields, music design may enable original research contributions in the fields of Engineering Science and the Humanities, but also in other fields such as Health and Natural Science. Also, music design may pave the way for a diversity of new music practices that will benefit from the situated characteristics of design acts. Eventually, music design may be a suitable practice to assist computer science students in the exploration of more sensible approaches to music research. As a conclusive statement, I would like to emphasize that this post-scriptum was not intended to discredit other research approaches to music—especially the applied science approach, which was at the core of the present dissertation. Rather, it was meant to make my perception of the current landscape audible, in the hope of setting a common ground for music research. I hope to hear back from other people about it.

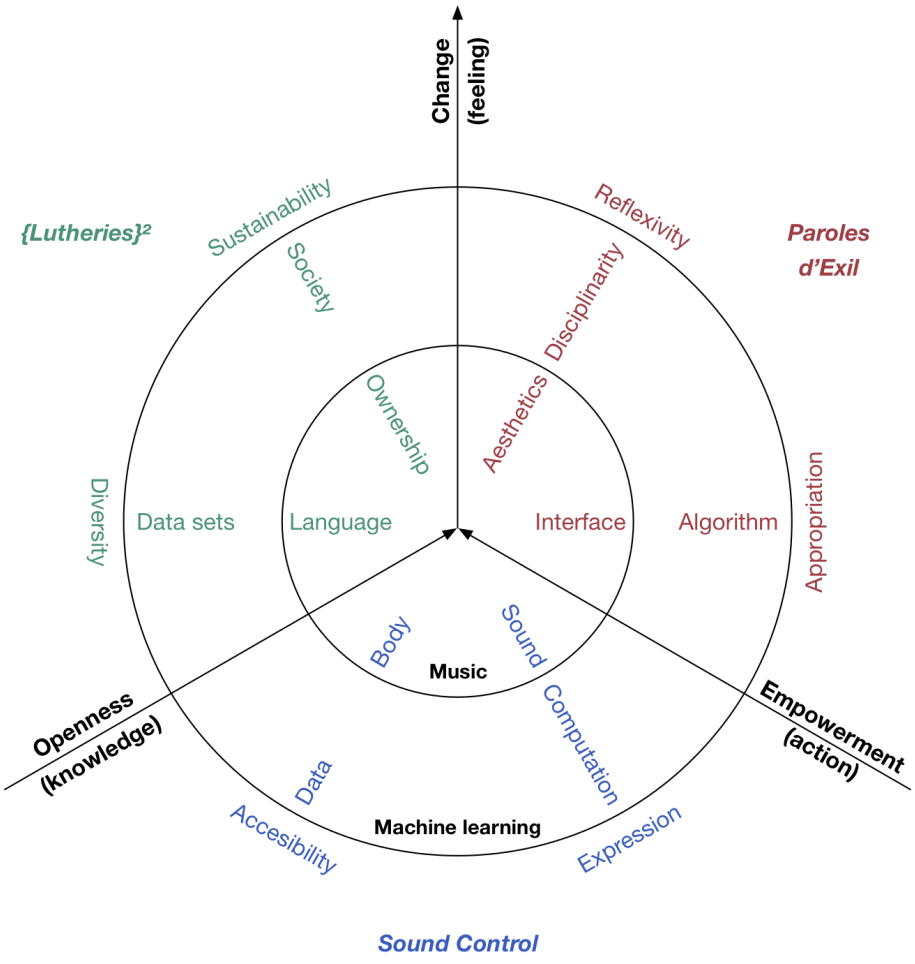


Figure I: The three example projects used design practice to shift the normative elements of music and machine learning toward openness, empowerment, and change.

Publications

As First Author

Scurto, H., & Chemla—Romeu-Santos, A. (2019). **Machine Learning for Computer Music Multidisciplinary Research: A Practical Case Study**. To appear in *International Symposium on Computer Music Multidisciplinary Research (CMMR'19).* Marseille, France. Springer, Cham.

Scurto, H., & Chemla—Romeu-Santos, A. (2019). **ægo**. To appear in *International Symposium on Computer Music Multidisciplinary Research (CMMR'19).* Marseille, France. Springer, Cham.

Scurto, H., Liu, W., Matuszewski, B., Bevilacqua, F., Fréchin, J.-L., Petrevski, U., & Schnell, N. (2019). **Entrain: Encouraging Social Interaction in Collective Music Making**. *ACM SIGGRAPH Studio 2019 (SIGGRAPH'19).* Los Angeles, California: ACM, 2019, pp.5—6.

Scurto, H., Van Kerrebroeck, B., Caramiaux, B., & Bevilacqua, F. (2019). **Designing Deep Reinforcement Learning for Human Parameter Exploration**. *arXiv preprint arXiv:1907.00824.*

Scurto, H., Van Kerrebroeck, B., & Bevilacqua, F., **Riding the Co-Explorers**. In *Proceedings of the 19th International Society for Music Information Retrieval Conference (ISMIR'18).* Paris, France, 2018.

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'18).* Limassol, Cyprus, 2018.

Scurto, H., & Bevilacqua, F., **Appropriating Music Computing Practices Through Human-AI Collaboration**. In *Actes des Journées d'Informatique Musicale (JIM'18).* Amiens, France, 2018.

Scurto, H., Bevilacqua, F., & Françoise, J., **Shaping and Exploring Interactive Motion-Sound Mappings Using Online Clustering Techniques**. In *Proceedings of the 17th International Conference on New Interfaces for Musical Expression (NIME'17).* Copenhagen, Denmark, 2017.

Scurto, H., & Bevilacqua, F., **VIMOs: Enabling Expressive Mediation and Generation of Embodied Musical Interactions**. In *Proceedings of the 4th International Conference on Movement Computing (MOCO'17)*. London, United Kingdom, 2017.

Scurto, H., & Fiebrink, R., **Grab-and-Play Mapping: Creative Machine Learning Approaches for Musical Inclusion and Exploration**. In *Proceedings of the 42nd International Computer Music Conference (ICMC'16)*. Utrecht, Netherlands, 2016.

Scurto, H., **Creative Interaction Paradigms for Machine Learning-Based Design of Gestural Musical Interfaces**. Pre-doctoral report, ENS Paris-Saclay, Goldsmiths University of London, 2016.

Scurto, H., Lemaitre, G., Françoise, J., Voisin, F., Bevilacqua, F., & Susini, P., **Embodying sounds: Building and analysis of a database of gestural and vocal imitations**. In *7th Conference of the International Society for Gesture Studies (ISGS'16)*. Paris, France, 2016.

Scurto, H., Lemaitre, G., Françoise, J., Voisin, F., Bevilacqua, F., & Susini, P., **Combining gestures and vocalizations to imitate sounds**. In *170th Meeting of the Acoustical Society of America (ASA'15)*. Jacksonville, USA, 2015.

Scurto, H., **Combination of Gesture and Vocalization in Sound Imitation**. Master's thesis, UPMC, IRCAM, 2015.

As Secondary Author

Parke-Wolfe, S., Scurto, H., & Fiebrink, R., **Sound Control: Supporting Custom Musical Interface Design for Children with Disabilities**. In *Proceedings of the 19th International Conference on New Interfaces for Musical Expression (NIME'19)*. Porto Alegre, Brazil, 2019.

Van Kerrebroeck, B., **Deep Reinforcement Learning for the Design of Musical Interaction**. Master's thesis, Sorbonne Université, IRCAM, 2018.

Lemaitre, G., Scurto, H., Françoise, J., Bevilacqua, F., Houix, O., & Susini, P, **Rising tones and rustling noises: Metaphors in gestural depictions of sounds**. *PloS one, 12(7), e0181786*. 2017.

Videography

Chapter 3

Scurto, H. (2017). *Online Gaussian Mixture Model: Interaction overview*. https://vimeo.com/300812404. (Cited on 32.)

Chapter 5

Scurto, H. (2018). *Co-Explorers: Interaction overview*. https://vimeo.com/300810165. (Cited on 77.)

Scurto, H. (2018). *Riding the Co-Explorers: Level overview*. https://youtu.be/ToAz3Zwssm0. (Cited on 89.)

Ecole des Arts Déco - Paris (2019). *Behavioral Matter*. https://vimeo.com/342930653. (Cited on 93.)

Chemla–Romeu-Santos, A., Scurto, H. (2019). *ægo: Improvisation overview*. https://youtu.be/gCzOoNChlJQ. (Cited on 101.)

Chapter 6

Scurto, H., *entrain: Interaction overview* (2019). https://vimeo.com/331404969. (Cited on 114.)

Bibliography

[Abbeel and Ng, 2004] Abbeel, P. and Ng, A. Y. (2004). Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM. (Cited on 88.)

[Agamben, 2009] Agamben, G. (2009). *" What is an apparatus?" and other essays*. Stanford University Press. (Cited on 146.)

[Alpern, 1995] Alpern, A. (1995). Techniques for algorithmic composition of music. *On the Web*. (Cited on 15.)

[Altavilla et al., 2013] Altavilla, A., Caramiaux, B., and Tanaka, A. (2013). Towards gestural sonic affordances. (Cited on 131.)

[Altenmüller et al., 2013] Altenmüller, E., Schmidt, S., and Zimmermann, E. (2013). *The evolution of emotional communication: From sounds in nonhuman mammals to speech and music in man*. OUP Oxford. (Cited on 13, 135, 157.)

[Alvarez et al., 2009] Alvarez, M., Luengo, D., and Lawrence, N. (2009). Latent force models. In *Artificial Intelligence and Statistics*, pages 9–16. (Cited on 18.)

[Amershi et al., 2014] Amershi, S., Cakmak, M., Knox, W. B., and Kulesza, T. (2014). Power to the people: The role of humans in interactive machine learning. *AI Magazine*, 35(4):105–120. (Cited on 22, 60, 63, 64, 80, 87, 128.)

[Amershi et al., 2015] Amershi, S., Chickering, M., Drucker, S. M., Lee, B., Simard, P., and Suh, J. (2015). Modeltracker: Redesigning performance analysis tools for machine learning. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 337–346. ACM. (Cited on 63.)

[Amershi et al., 2012] Amershi, S., Fogarty, J., and Weld, D. (2012). Regroup: Interactive machine learning for on-demand group creation in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 21–30. ACM. (Cited on 63.)

[Amershi et al., 2011] Amershi, S., Lee, B., Kapoor, A., Mahajan, R., and Christian, B. (2011). Cuet: human-guided fast and accurate network alarm triage. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 157–166. ACM. (Cited on 63.)

[Amershi et al., 2019] Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., et al. (2019). Guidelines for human-ai interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, page 3. ACM. (Cited on 84.)

[Andersen and Knees, 2016] Andersen, K. and Knees, P. (2016). Conversations with expert users in music retrieval and research challenges for creative mir. In *ISMIR*, pages 122–128. (Cited on 88, 123.)

[Archer, 1979] Archer, B. (1979). Design as a discipline. *Design studies*, 1(1):17–20. (Cited on 145.)

[Assayag, 2014] Assayag, G. (2014). Creative symbolic interaction. In *40th Intl. Comp. Mus. Conf. and 11th Sound and Music Comp. Conf.(ICMC/SMC joint conf.)*, pages pp–1. ICMA, SMC, National and Kapodistrian University of Athens, IRMA. (Cited on 45, 136.)

[Assayag et al., 2006] Assayag, G., Bloch, G., Chemillier, M., Cont, A., and Dubnov, S. (2006). Omax brothers: a dynamic topology of agents for improvization learning. In *Proceedings of the 1st ACM workshop on Audio and music computing multimedia*, pages 125–132. ACM. (Cited on 1, 20, 43, 45, 99.)

[Assayag et al., 1999] Assayag, G., Rueda, C., Laurson, M., Agon, C., and Delerue, O. (1999). Computer-assisted composition at ircam: From patchwork to openmusic. *Computer Music Journal*, 23(3):59–72. (Cited on 15.)

[Athukorala et al., 2016a] Athukorala, K., Medlar, A., Oulasvirta, A., Jacucci, G., and Glowacka, D. (2016a). Beyond relevance: Adapting exploration/exploitation in information retrieval. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*, pages 359–369. ACM. (Cited on 60.)

[Athukorala et al., 2016b] Athukorala, K., Medlar, A., Oulasvirta, A., Jacucci, G., and Glowacka, D. (2016b). Beyond relevance: Adapting exploration/exploitation in information retrieval. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*, IUI '16, pages 359–369, New York, NY, USA. ACM. (Cited on 64.)

[Auger, 2013] Auger, J. (2013). Speculative design: crafting the speculation. *Digital Creativity*, 24(1):11–35. (Cited on 145.)

[Ballet et al., 1999] Ballet, G., Borghesi, R., Hoffmann, P., and Lévy, F. (1999). Studio online 3.0: An internet" killer application" for remote access to ircam sounds and processing tools. *Journée d'Informatique Musicale (JIM)*. (Cited on 103.)

[Bardzell and Bardzell, 2013] Bardzell, J. and Bardzell, S. (2013). What is critical about critical design? In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 3297–3306. ACM. (Cited on 145.)

[Beaudouin-Lafon, 2004] Beaudouin-Lafon, M. (2004). Designing interaction, not interfaces. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, AVI '04, pages 15–22, New York, NY, USA. ACM. (Cited on 14, 19, 22, 125.)

[Beaudouin-Lafon and Mackay, 2009] Beaudouin-Lafon, M. and Mackay, W. E. (2009). Prototyping tools and techniques. In *Human-Computer Interaction*, pages 137–160. CRC Press. (Cited on 4, 24, 59, 125, 140.)

[Beaudouin-Lafon and Mackay, 2018] Beaudouin-Lafon, M. and Mackay, W. E. (2018). Rethinking interaction: From instrumental interaction to human-computer partnerships. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, page W34. ACM. (Cited on 2.)

[Bell and Matuszewski, 2017] Bell, J. and Matuszewski, B. (2017). Smartvox–a web-based distributed media player as notation tool for choral practices. In *Proceedings of the International Conference on Technologies for Music Notation and Representation–TENOR*. (Cited on 16.)

[Bellemare et al., 2016] Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. (2016). Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pages 1471–1479. (Cited on 75.)

[Bertin-Mahieux et al., 2011] Bertin-Mahieux, T., Ellis, D. P., Whitman, B., and Lamere, P. (2011). The million song dataset. In *ISMIR 2011: Proceedings of the 12th International Society for Music Information Retrieval Conference, October 24-28, 2011, Miami, Florida*, pages 591–596. University of Miami. (Cited on 17, 148.)

[Bevilacqua et al., 2007] Bevilacqua, F., Guédy, F., Schnell, N., Fléty, E., and Leroy, N. (2007). Wireless sensor interface and gesture-follower for music pedagogy. In *Proceedings of the 7th International Conference on New Interfaces for Musical Expression (NIME)*, pages 124–129. ACM. (Cited on 44.)

[Bevilacqua et al., 2005] Bevilacqua, F., Müller, R., and Schnell, N. (2005). Mnm: a max/msp mapping toolbox. In *Proceedings of the 2005 conference on New interfaces for musical expression*, pages 85–88. National University of Singapore. (Cited on 23.)

[Bevilacqua and Schnell, 2016] Bevilacqua, F. and Schnell, N. (2016). From musical interfaces to musical interactions, in human computer confluence: Transforming human experience through symbiotic technologies. (Cited on 14.)

[Bevilacqua et al., 2009] Bevilacqua, F., Zamborlin, B., Sypniewski, A., Schnell, N., Guédy, F., and Rasamimanana, N. (2009). Continuous realtime gesture following and recognition. In *International gesture workshop*, pages 73–84. Springer. (Cited on 1, 25, 28, 29, 105.)

[Bianchini and Verhagen, 2016] Bianchini, S. and Verhagen, E. (2016). *Practicable: From participation to interaction in contemporary art*. MIT Press. (Cited on 2, 6, 140, 146.)

[Bishop, 2006] Bishop, C. M. (2006). *Pattern recognition and machine learning*. springer. (Cited on 17, 18, 125, 134.)

[Blythe et al., 2016] Blythe, M., Andersen, K., Clarke, R., and Wright, P. (2016). Anti-solutionist strategies: Seriously silly design fiction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4968–4978. ACM. (Cited on 88.)

[Bode, 1984] Bode, H. (1984). History of electronic sound modification. *Journal of the Audio Engineering Society*, 32(10):730–739. (Cited on 13.)

[Born, 1995] Born, G. (1995). *Rationalizing culture: IRCAM, Boulez, and the institutionalization of the musical avant-garde*. Univ of California Press. (Cited on 12, 148.)

[Boulanger-Lewandowski et al., 2012] Boulanger-Lewandowski, N., Bengio, Y., and Vincent, P. (2012). Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. *arXiv preprint arXiv:1206.6392*. (Cited on 19.)

[Bourganel et al., 2016] Bourganel, R., Hugon, S., and Badr, M. (2016). Algorithmes, formes diachroniques et métadesign. *Sciences du Design*, 4:98–111. (Cited on 20, 152.)

[Bresson et al., 2011] Bresson, J., Agon, C., and Assayag, G. (2011). Openmusic: visual programming environment for music composition, analysis and research. In *Proceedings of the 19th ACM international conference on Multimedia*, pages 743–746. ACM. (Cited on 123.)

[Briot et al., 2017] Briot, J.-P., Hadjeres, G., and Pachet, F. (2017). Deep learning techniques for music generation-a survey. *arXiv preprint arXiv:1709.01620*. (Cited on 1, 21, 126.)

[Brockman et al., 2016] Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). Openai gym. *arXiv preprint arXiv:1606.01540*. (Cited on 88.)

[Buchanan, 2005] Buchanan, B. G. (2005). A (very) brief history of artificial intelligence. *Ai Magazine*, 26(4):53–53. (Cited on 124.)

[Burnham, 1968] Burnham, J. (1968). Systems esthetics. *Artforum*, 7(1):30–35. (Cited on 146.)

[Cadoz and Wanderley, 2000] Cadoz, C. and Wanderley, M. M. (2000). Gesture-music. (Cited on 12.)

[Cakmak et al., 2010] Cakmak, M., Chao, C., and Thomaz, A. L. (2010). Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118. (Cited on 23.)

[Candy, 2006] Candy, L. (2006). Practice based research: A guide. *CCS Report*, 1:1–19. (Cited on 145.)

[Canonne, 2012] Canonne, C. (2012). Improvisation collective libre et processus de création musicale: création et créativité au prisme de la coordination. *Rev. Musicol*, 98:107–148. (Cited on 135.)

[Caramiaux et al., 2014a] Caramiaux, B., Bevilacqua, F., Bianco, T., Schnell, N., Houix, O., and Susini, P. (2014a). The role of sound source perception in gestural sound description. *ACM Transactions on Applied Perception (TAP)*, 11(1):1. (Cited on 13.)

[Caramiaux et al., 2014b] Caramiaux, B., Françoise, J., Schnell, N., and Bevilacqua, F. (2014b). Mapping through listening. *Computer Music Journal*, 38(3):34–48. (Cited on 131.)

[Caramiaux et al., 2019] Caramiaux, B., Lotte, F., Geurts, J., Amato, G., Behrmann, M., Bimbot, F., Falchi, F., Garcia, A., Gibert, J., Gravier, G., et al. (2019). Ai in the media and creative industries. *arXiv preprint arXiv:1905.04175*. (Cited on 84, 125, 150.)

[Caramiaux et al., 2015] Caramiaux, B., Montecchio, N., Tanaka, A., and Bevilacqua, F. (2015). Adaptive gesture recognition with variation estimation for interactive systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 4(4):18. (Cited on 25, 29.)

[Caramiaux and Tanaka, 2013] Caramiaux, B. and Tanaka, A. (2013). Machine learning of musical gestures. In *NIME*, pages 513–518. (Cited on 29, 45.)

[Carpentier et al., 2015] Carpentier, T., Noisternig, M., and Warusfel, O. (2015). Twenty years of ircam spat: looking back, looking forward. In *41st International Computer Music Conference (ICMC)*, pages 270–277. (Cited on 43.)

[Cartwright and Pardo, 2014] Cartwright, M. and Pardo, B. (2014). Synthassist: Querying an audio synthesizer by vocal imitation. In *Proceedings of the 14th International Conference on New Interfaces for Musical Expression (NIME)*. (Cited on 43.)

[Cartwright and Pardo, 2015] Cartwright, M. and Pardo, B. (2015). Vocalsketch: Vocally imitating audio concepts. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 43–46. ACM. (Cited on 15.)

[Cartwright et al., 2014] Cartwright, M., Pardo, B., and Reiss, J. (2014). Mixploration: Rethinking the audio mixer interface. In *Proceedings of the 19th international conference on Intelligent User Interfaces*, pages 365–370. ACM. (Cited on 43, 62.)

[Cartwright and Pardo, 2013] Cartwright, M. B. and Pardo, B. (2013). Social-eq: Crowdsourcing an equalization descriptor map. In *ISMIR*, pages 395–400. (Cited on 16.)

[Cath et al., 2018] Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., and Floridi, L. (2018). Artificial intelligence and the 'good society': the us, eu, and uk approach. *Science and engineering ethics*, 24(2):505–528. (Cited on 150.)

[Chadabe, 1977] Chadabe, J. (1977). Some reflections on the nature of the landscape. *Computer Music Journal*, 1(3):5–11. (Cited on 15.)

[Chemla et al., 2015] Chemla, A. et al. (2015). Guidages de l'improvisation. Master's thesis, Ircam UMR STMS 9912. (Cited on 39.)

[Cherry and Latulipe, 2014] Cherry, E. and Latulipe, C. (2014). Quantifying the creativity support of digital tools through the creativity support index. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 21(4):21. (Cited on 85.)

[Chowning, 1973] Chowning, J. M. (1973). The synthesis of complex audio spectra by means of frequency modulation. *Journal of the audio engineering society*, 21(7):526–534. (Cited on 50, 67, 104.)

[Christiano et al., 2017] Christiano, P., Leike, J., Brown, T. B., Martic, M., Legg, S., and Amodei, D. (2017). Deep reinforcement learning from human preferences. *arXiv preprint arXiv:1706.03741*. (Cited on 46, 48, 57, 64.)

[Cockburn et al., 2015] Cockburn, A., Gutwin, C., Scarr, J., and Malacria, S. (2015). Supporting novice to expert transitions in user interfaces. *ACM Computing Surveys (CSUR)*, 47(2):31. (Cited on 43, 44.)

[Colton et al., 2012] Colton, S., Wiggins, G. A., et al. (2012). Computational creativity: The final frontier? In *Ecai*, volume 12, pages 21–26. Montpelier. (Cited on 21.)

[Consolvo et al., 2009] Consolvo, S., McDonald, D. W., and Landay, J. A. (2009). Theory-driven design strategies for technologies that support behavior change in everyday life. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 405–414. ACM. (Cited on 24.)

[Cont, 2008] Cont, A. (2008). Antescofo: Anticipatory synchronization and control of interactive parameters in computer music. In *International Computer Music Conference (ICMC)*, pages 33–40. (Cited on 44.)

[Crandall et al., 2018] Crandall, J. W., Oudah, M., Ishowo-Oloko, F., Abdallah, S., Bonnefon, J.-F., Cebrian, M., Shariff, A., Goodrich, M. A., Rahwan, I., et al. (2018). Cooperating with machines. *Nature communications*, 9(1):233. (Cited on 64, 88.)

[Crawford and Joler, 2018] Crawford, K. and Joler, V. (2018). Anatomy of an ai system. *Retrieved September*, 18:2018. (Cited on 150.)

[Cross, 2001a] Cross, I. (2001a). Music, cognition, culture, and evolution. *Annals of the New York Academy of sciences*, 930(1):28–42. (Cited on 21.)

[Cross, 2001b] Cross, N. (2001b). Designerly ways of knowing: Design discipline versus design science. *Design issues*, 17(3):49–55. (Cited on 145.)

[Csikszentmihalyi, 1997] Csikszentmihalyi, M. (1997). Flow and the psychology of discovery and invention. *HarperPerennial, New York*, 39. (Cited on 40, 87, 124.)

[Davis et al., 2016] Davis, N., Hsiao, C.-P., Yashraj Singh, K., Li, L., and Magerko, B. (2016). Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*, pages 196–207. ACM. (Cited on 62, 86.)

[Davis et al., 2014] Davis, N. M., Popova, Y., Sysoev, I., Hsiao, C.-P., Zhang, D., and Magerko, B. (2014). Building artistic computer colleagues with an enactive model of creativity. In *ICCC*, pages 38–45. (Cited on 62.)

[Deleuze, 1995] Deleuze, G. (1995). Post-script on control societies. *Les éditions de Minuit*. (Cited on 150.)

[Delle Monache et al., 2018] Delle Monache, S., Rocchesso, D., Bevilacqua, F., Lemaitre, G., Baldan, S., and Cera, A. (2018). Embodied sound design. *International Journal of Human-Computer Studies*. (Cited on 44, 61.)

[Derbinsky and Essl, 2012] Derbinsky, N. and Essl, G. (2012). Exploring reinforcement learning for mobile percussive collaboration. In *NIME*. (Cited on 47.)

[Deruty, 2016] Deruty, E. (2016). Goal-oriented mixing. In *Proceedings of the 2nd AES Workshop on Intelligent Music Production*, volume 13. (Cited on 16, 19, 125.)

[Deterding et al., 2017] Deterding, C. S., Hook, J. D., Fiebrink, R., Gow, J., Akten, M., Smith, G., Liapis, A., and Compton, K. (2017). Mixed-initiative creative interfaces. In *CHI EA'17: Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM. (Cited on 63.)

[d'Inverno and McCormack, 2015] d'Inverno, M. and McCormack, J. (2015). Heroic versus collaborative ai for the arts. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*. (Cited on 88.)

[Dix, 2007] Dix, A. (2007). Designing for appropriation. In *Proceedings of the 21st British HCI Group Annual Conference on People and Computers: HCI... but not as we know it-Volume 2*, pages 27–30. British Computer Society. (Cited on 62.)

[Dobrian and Koppelman, 2006] Dobrian, C. and Koppelman, D. (2006). The'e'in nime: musical expression with new computer interfaces. In *Proceedings of the 2006 conference on New interfaces for musical expression*, pages 277–282. IRCAM—Centre Pompidou. (Cited on 130.)

[Domingos, 2012] Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10):78–87. (Cited on 22, 149.)

[Donin, 2004] Donin, N. (2004). Towards organised listening: some aspects of the 'signed listening'project, ircam. *Organised Sound*, 9(1):99–108. (Cited on 12.)

[Dorst and Cross, 2001] Dorst, K. and Cross, N. (2001). Creativity in the design process: co-evolution of problem–solution. *Design studies*, 22(5):425–437. (Cited on 61, 63.)

[Dourish, 2004] Dourish, P. (2004). *Where the action is: the foundations of embodied interaction*. MIT press. (Cited on 22, 131.)

[Dove et al., 2017] Dove, G., Halskov, K., Forlizzi, J., and Zimmerman, J. (2017). 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. ACM. (Cited on 62.)

[Downie, 2003] Downie, J. S. (2003). Music information retrieval. *Annual review of information science and technology*, 37(1):295–340. (Cited on 19.)

[Drucker, 2011] Drucker, J. (2011). Humanities approaches to graphical display. *Digital Humanities Quarterly*, 5(1):1–21. (Cited on 148.)

[Dryhurst, 2019] Dryhurst, M. (2019). Band together: why musicians must strike a collective chord to survive. *The Guardian*. (Cited on 148.)

[Dubnov et al., 2003] Dubnov, S., Assayag, G., Lartillot, O., and Bejerano, G. (2003). Using machine-learning methods for musical style modeling. *Computer*, 36(10):73–80. (Cited on 19.)

[Dubnov and Surges, 2014] Dubnov, S. and Surges, G. (2014). Delegating creativity: Use of musical algorithms in machine listening and composition. In *Digital Da Vinci*, pages 127–158. Springer. (Cited on 123.)

[Eigenfeldt and Pasquier, 2013] Eigenfeldt, A. and Pasquier, P. (2013). Evolving structures for electronic dance music. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 319–326. ACM. (Cited on 21.)

[Einbond et al., 2016] Einbond, A., Schwarz, D., Borghesi, R., and Schnell, N. (2016). Introducing catoracle: Corpus-based concatenative improvisation with the audio oracle algorithm. In *Proceedings of the International Computer Music Conference*. (Cited on 153.)

[Esling et al., 2018] Esling, P., Chemla-Romeu-Santos, A., and Bitton, A. (2018). Bridging audio analysis, perception and synthesis with perceptually-regularized variational timbre spaces. In *ISMIR*, pages 175–181. (Cited on 19, 98.)

[Esling et al., 2019] Esling, P., Masuda, N., Bardet, A., Despres, R., et al. (2019). Universal audio synthesizer control with normalizing flows. *arXiv preprint arXiv:1907.00971*. (Cited on 1, 136.)

[Fails and Olsen Jr, 2003] Fails, J. A. and Olsen Jr, D. R. (2003). Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces*, pages 39–45. ACM. (Cited on 23, 63.)

[Fallman, 2003] Fallman, D. (2003). Design-oriented human-computer interaction. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 225–232. ACM. (Cited on 27.)

[Fels and Hinton, 1993] Fels, S. S. and Hinton, G. E. (1993). Glove-talk: A neural network interface between a data-glove and a speech synthesizer. *IEEE transactions on Neural Networks*, 4(1):2–8. (Cited on 23.)

[Fiebrink, 2017] Fiebrink, R. (2017). Machine learning as meta-instrument: Human-machine partnerships shaping expressive instrumental creation. In *Musical Instruments in the 21st Century*, pages 137–151. Springer. (Cited on 25.)

[Fiebrink, 2019] Fiebrink, R. (2019). Machine learning education for artists, musicians, and other creative practitioners. *ACM Transactions on Computing Education*. (Cited on 23, 25, 87.)

[Fiebrink and Caramiaux, 2016] Fiebrink, R. and Caramiaux, B. (2016). The machine learning algorithm as creative musical tool. *Handbook of Algorithmic Music*. (Cited on 45, 47, 63, 88, 99, 126.)

[Fiebrink et al., 2011] Fiebrink, R., Cook, P. R., and Trueman, D. (2011). Human model evaluation in interactive supervised learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 147–156, New York, NY, USA. ACM. (Cited on 23, 25, 29, 58, 63, 85, 86, 87, 139.)

[Fiebrink et al., 2010] Fiebrink, R., Trueman, D., Britt, N. C., Nagai, M., Kaczmarek, K., Early, M., Daniel, M., Hege, A., and Cook, P. R. (2010). Toward understanding human-computer interaction in composing the instrument. In *ICMC*. (Cited on 23, 25, 28, 29, 45, 57, 63.)

[Fiebrink et al., 2009] Fiebrink, R., Trueman, D., and Cook, P. R. (2009). A meta-instrument for interactive, on-the-fly machine learning. In *NIME*, pages 280–285. (Cited on 24.)

[Fiebrink, 2011] Fiebrink, R. A. (2011). *Real-time human interaction with supervised learning algorithms for music composition and performance*. PhD thesis, Princeton University. (Cited on 2, 23.)

[Findeli, 2004] Findeli, A. (2004). La recherche-projet: une méthode pour la recherche en design. In *Texte de la communication présentée au premier Symposium de recherche sur le design à HGK de Bâle sous les auspices du Swiss Design Network, Bâle*. (Cited on 145.)

[Fitzgerald et al., 2017] Fitzgerald, T., Goel, A. K., and Thomaz, A. (2017). Human-robot co-creativity: Task transfer on a spectrum of similarity. In *ICCC*, pages 104–111. (Cited on 64.)

[Fogel, 2006] Fogel, D. B. (2006). *Evolutionary computation: toward a new philosophy of machine intelligence*, volume 1. John Wiley & Sons. (Cited on 88.)

[Françoise, 2015] Françoise, J. (2015). *Motion-sound mapping by demonstration*. PhD thesis, UPMC. (Cited on 1, 2.)

[Francoise and Bevilacqua, 2018] Francoise, J. and Bevilacqua, F. (2018). Motion-sound mapping through interaction: An approach to user-centered design of auditory feedback using machine learning. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 8(2):16. (Cited on 63.)

[Françoise et al., 2014] Françoise, J., Schnell, N., Borghesi, R., and Bevilacqua, F. (2014). Probabilistic models for designing motion and sound relationships. In *Proceedings of the 2014 international conference on new interfaces for musical expression*, pages 287–292. (Cited on 25, 28, 29, 30, 33, 45.)

[Franinović and Serafin, 2013] Franinović, K. and Serafin, S. (2013). *Sonic interaction design*. Mit Press. (Cited on 156.)

[Frenoy et al., 2016] Frenoy, R., Soullard, Y., Thouvenin, I., and Gapenne, O. (2016). Adaptive training environment without prior knowledge: Modeling feedback selection as a multi-armed bandit problem. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization*, pages 131–139. ACM. (Cited on 64.)

[Frich et al., 2018] Frich, J., Biskjaer, M. M., and Dalsgaard, P. (2018). Twenty years of creativity research in human-computer interaction: Current state and future directions. In *Designing Interactive Systems*. Association for Computing Machinery (ACM). (Cited on 22, 62, 135.)

[Frith, 1998] Frith, S. (1998). *Performing rites: On the value of popular music*. Harvard University Press. (Cited on 12, 125, 156.)

[Garcia, 2014] Garcia, J. (2014). *Supporting music composition with interactive paper*. PhD thesis, Université Paris Sud-Paris XI. (Cited on 22, 121.)

[Garcia et al., 2012] Garcia, J., Tsandilas, T., Agon, C., and Mackay, W. (2012). Interactive paper substrates to support musical creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1825–1828. ACM. (Cited on 3, 61.)

[Gaver, 1991] Gaver, W. W. (1991). Technology affordances. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 79–84. Acm. (Cited on 131.)

[Gentes, 2017] Gentes, A. (2017). *The In-Discipline of Design: Bridging the Gap Between Humanities and Engineering*. Springer. (Cited on 145.)

[Geraci, 2012] Geraci, R. M. (2012). *Apocalyptic AI: Visions of heaven in robotics, artificial intelligence, and virtual reality*. Oxford University Press. (Cited on 124.)

[Ghisi, 2017] Ghisi, D. (2017). *Music across music: towards a corpus-based, interactive computer-aided composition*. PhD thesis, Paris 6. (Cited on 141.)

[Ghomi, 2012] Ghomi, E. (2012). *Designing expressive interaction techniques for novices inspired by expert activities: the case of musical practice*. PhD thesis, Université Paris Sud-Paris XI. (Cited on 22.)

[Giaccardi and Fischer, 2008] Giaccardi, E. and Fischer, G. (2008). Creativity and evolution: a metadesign perspective. *Digital Creativity*, 19(1):19–32. (Cited on 135, 149, 156.)

[Gibson, 1977] Gibson, J. J. (1977). The theory of affordances. *Hilldale, USA*, 1(2). (Cited on 131, 147.)

[Gillies, 2019] Gillies, M. (2019). Understanding the role of interactive machine learning in movement interaction design. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(1):5. (Cited on 63.)

[Gillies et al., 2016] Gillies, M., Fiebrink, R., Tanaka, A., Garcia, J., Bevilacqua, F., Heloir, A., Nunnari, F., Mackay, W., Amershi, S., Lee, B., et al. (2016). Human-centred machine learning. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 3558–3565. ACM. (Cited on 2, 3, 63, 148.)

[Glowacka et al., 2013] Glowacka, D., Ruotsalo, T., Konuyshkova, K., Athukorala, k., Kaski, S., and Jacucci, G. (2013). Directing exploratory search: Reinforcement learning from user interactions with keywords. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces*, IUI '13, pages 117–128, New York, NY, USA. ACM. (Cited on 64.)

[Godøy and Leman, 2010] Godøy, R. I. and Leman, M. (2010). *Musical gestures: Sound, movement, and meaning*. Routledge. (Cited on 12.)

[Goodfellow et al., 2016] Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT press. (Cited on 18.)

[Hannon and Trainor, 2007] Hannon, E. E. and Trainor, L. J. (2007). Music acquisition: effects of enculturation and formal training on development. *Trends in cognitive sciences*, 11(11):466–472. (Cited on 13.)

[Hart et al., 2017] Hart, Y., Mayo, A. E., Mayo, R., Rozenkrantz, L., Tendler, A., Alon, U., and Noy, L. (2017). Creative foraging: An experimental paradigm for studying exploration and discovery. *PloS one*, 12(8):e0182133. (Cited on 61, 135.)

[Haynes, 2003] Haynes, R. (2003). From alchemy to artificial intelligence: Stereotypes of the scientist in western literature. *Public Understanding of Science*, 12(3):243–253. (Cited on 124, 132.)

[Henderson, 1983] Henderson, S. M. (1983). Effects of a music therapy program upon awareness of mood in music, group cohesion, and self-esteem among hospitalized adolescent patients. *Journal of music therapy*, 20(1):14–20. (Cited on 13.)

[Herbert, 2012] Herbert, T. (2012). Social history and music history. In *The Cultural Study of Music*, pages 71–80. Routledge. (Cited on 14, 156.)

[Hiller and Isaacson, 1957] Hiller, L. and Isaacson, L. M. (1957). *Illiac suite, for string quartet*, volume 30. New Music Edition. (Cited on 15.)

[Hiller Jr and Isaacson, 1958] Hiller Jr, L. A. and Isaacson, L. M. (1958). Musical composition with a high-speed digital computer. *Journal of the Audio Engineering Society*, 6(3):154–160. (Cited on 123.)

[Höök and Löwgren, 2012] Höök, K. and Löwgren, J. (2012). Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 19(3):23. (Cited on 4, 40, 133.)

[Horvitz, 1999] Horvitz, E. (1999). Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 159–166. ACM. (Cited on 63, 84.)

[Huang et al., 2019] Huang, C.-Z. A., Hawthorne, C., Roberts, A., Dinculescu, M., Wexler, J., Hong, L., and Howcroft, J. (2019). The bach doodle: Approachable music composition with machine learning at scale. *arXiv preprint arXiv:1907.06637*. (Cited on 19.)

[Hunt and Wanderley, 2002] Hunt, A. and Wanderley, M. M. (2002). Mapping performer parameters to synthesis engines. *Organised sound*, 7(2):97–108. (Cited on 15, 28, 29.)

[Hutchinson et al., 2003] Hutchinson, H., Mackay, W., Westerlund, B., Bederson, B. B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H., et al. (2003). Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 17–24. ACM. (Cited on 24, 84.)

[Hyde, 2010] Hyde, L. (2010). *Common as air: Revolution, art, and ownership*. Farrar, Straus and Giroux. (Cited on 148.)

[Illari et al., 2011] Illari, P. M., Russo, F., and Williamson, J. (2011). *Causality in the Sciences*. Oxford University Press. (Cited on 149.)

[Jolliffe, 2011] Jolliffe, I. (2011). Principal component analysis. In *International encyclopedia of statistical science*, pages 1094–1096. Springer. (Cited on 81.)

[Jorda, 2005] Jorda, S. (2005). *Digital Lutherie Crafting musical computers for new musics' performance and improvisation*. PhD thesis, Universitat Pompeu Fabra. (Cited on 1, 2, 13, 14, 15, 16, 30, 33, 57, 68, 127, 130, 132.)

[Jost, 2013] Jost, C. (2013). The performative turn in music research between desideratum and (partially) disciplinary paradigm christofer jost. *MUSIKTHEORIE*, 28(4):291–309. (Cited on 12.)

[Kantosalo et al., 2014] Kantosalo, A., Toivanen, J. M., Xiao, P., and Toivonen, H. (2014). From isolation to involvement: Adapting machine creativity software to support human-computer co-creation. In *ICCC*, pages 1–7. (Cited on 62.)

[Kapoor et al., 2010] Kapoor, A., Lee, B., Tan, D., and Horvitz, E. (2010). Interactive optimization for steering machine classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1343–1352. (Cited on 23, 63.)

[Katan et al., 2015] Katan, S., Grierson, M., and Fiebrink, R. (2015). 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. ACM. (Cited on 63.)

[Kingma and Welling, 2013] Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*. (Cited on 98.)

[Kläy et al., 2015] Kläy, A., Zimmermann, A. B., and Schneider, F. (2015). Rethinking science for sustainable development: Reflexive interaction for a paradigm transformation. *Futures*, 65:72–85. (Cited on 5.)

[Kleinsmith and Gillies, 2013] Kleinsmith, A. and Gillies, M. (2013). Customizing by doing for responsive video game characters. *International Journal of Human-Computer Studies*, 71(7-8):775–784. (Cited on 63.)

[Knox and Stone, 2009] Knox, W. B. and Stone, P. (2009). Interactively shaping agents via human reinforcement: The tamer framework. In *Proceedings of the fifth international conference on Knowledge capture*, pages 9–16. ACM. (Cited on 46, 47, 48, 57, 60, 64, 74, 76.)

[Koch, 2017] Koch, J. (2017). Design implications for designing with a collaborative ai. In *2017 AAAI Spring Symposium Series*. (Cited on 62.)

[Koch et al., 2019] Koch, J., Lucero, A., Hegemann, L., and Oulasvirta, A. (2019). May ai?: Design ideation with cooperative contextual bandits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, page 633. ACM. (Cited on 85, 86.)

[Koch and Oulasvirta, 2018] Koch, J. and Oulasvirta, A. (2018). Group cognition and collaborative ai. In *Human and Machine Learning*, pages 293–312. Springer. (Cited on 64.)

[Kronland-Martinet, 1988] Kronland-Martinet, R. (1988). The wavelet transform for analysis, synthesis, and processing of speech and music sounds. *Computer Music Journal*, 12(4):11–20. (Cited on 97.)

[Kulic et al., 2007] Kulic, D., Takano, W., and Nakamura, Y. (2007). Incremental learning of full body motions via adaptive factorial hidden markov models. In *Proc. of the International Conference on Epigenetic Robotics*. (Cited on 39, 40.)

[Kumar et al., 2011] Kumar, R., Talton, J. O., Ahmad, S., and Klemmer, S. R. (2011). Bricolage: example-based retargeting for web design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2197–2206. ACM. (Cited on 63.)

[Latour, 1987] Latour, B. (1987). *Science in action: How to follow scientists and engineers through society*. Harvard university press. (Cited on 21.)

[Latour, 1996] Latour, B. (1996). *Petite réflexion sur le culte moderne des dieux faitiches*. (Cited on 157.)

[Lazar et al., 2017] Lazar, J., Feng, J. H., and Hochheiser, H. (2017). *Research methods in human-computer interaction*. Morgan Kaufmann. (Cited on 24, 41.)

[Lee and Wessel, 1992] Lee, M. and Wessel, D. (1992). Connectionist models for real-time control of synthesis and compositional algorithms. In *Proceedings of the International Computer Music Conference*, pages 277–277. (Cited on 23.)

[Leipp, 1972] Leipp, E. (1972). Les ondes martenot. *bulletin du GAM no60.* (Cited on 14.)

[Lemaitre et al., 2017] Lemaitre, G., Scurto, H., Françoise, J., Bevilacqua, F., Houix, O., and Susini, P. (2017). Rising tones and rustling noises: Metaphors in gestural depictions of sounds. *PloS one*, 12(7):e0181786. (Cited on 13.)

[Leman, 2008] Leman, M. (2008). *Embodied music cognition and mediation technology.* Mit Press. (Cited on 12, 16, 30, 43, 108, 122, 131, 145, 156.)

[Leman, 2016] Leman, M. (2016). *The expressive moment: How interaction (with music) shapes human empowerment.* MIT press. (Cited on 1, 3, 12, 13, 39, 44, 108, 124, 129, 130, 131, 132, 134, 135, 147, 157.)

[Levitin, 2006] Levitin, D. J. (2006). *This is your brain on music: The science of a human obsession.* Penguin. (Cited on 132.)

[Lévy, 2013] Lévy, B. (2013). *Principes et architectures pour un système interactif et agnostique dédié à l'improvisation musicale.* PhD thesis, Paris 6. (Cited on 20, 21.)

[Li, 2018] Li, Y. (2018). Deep reinforcement learning. *arXiv preprint arXiv:1810.06339.* (Cited on 60.)

[Lindh, 2018] Lindh, M. (2018). Beyond a skeuomorphic representation of subtractive synthesis. In *IUI Workshops.* (Cited on 147.)

[Liu, 2018] Liu, W. (2018). *Information Theory as a Unified Tool for Understanding and Designing Human-Computer Interaction.* PhD thesis, Université Paris-Saclay. (Cited on 31.)

[Liu et al., 2017] Liu, W., d'Oliveira, R. L., Beaudouin-Lafon, M., and Rioul, O. (2017). Bignav: Bayesian information gain for guiding multiscale navigation. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 5869–5880. ACM. (Cited on 113.)

[Lomas et al., 2016] Lomas, J. D., Forlizzi, J., Poonwala, N., Patel, N., Shodhan, S., Patel, K., Koedinger, K., and Brunskill, E. (2016). Interface design optimization as a multi-armed bandit problem. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4142–4153. ACM. (Cited on 49, 64, 66, 85.)

[Loy, 1985] Loy, G. (1985). Musicians make a standard: the midi phenomenon. *Computer Music Journal*, 9(4):8–26. (Cited on 15.)

[Maaten and Hinton, 2008] Maaten, L. v. d. and Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605. (Cited on 88.)

[Mackay, 1990] Mackay, W. E. (1990). *Users and customizable software: A co-adaptive phenomenon.* PhD thesis, Citeseer. (Cited on 86, 132.)

[Mackay, 2000] Mackay, W. E. (2000). Responding to cognitive overload: Co-adaptation between users and technology. *Intellectica*, 30(1):177–193. (Cited on 47.)

[Mackay and Fayard, 1997] Mackay, W. E. and Fayard, A.-L. (1997). Hci, natural science and design: a framework for triangulation across disciplines. In *Symposium on Designing Interactive Systems: Proceedings of the 2nd conference on Designing interactive systems: processes, practices, methods, and techniques*, pages 223–234. (Cited on 4, 21, 24.)

[Magnusson, 2019] Magnusson, T. (2019). *Sonic writing: technologies of material, symbolic, and signal inscriptions*. Bloomsbury Academic. (Cited on 147.)

[Matuszewski et al., 2018] Matuszewski, B., Larralde, J., and Bevilacqua, F. (2018). Designing movement driven audio applications using a web-based interactive machine learning toolkit. In *Web Audio Conference (WAC)*. (Cited on 25.)

[Matuszewski et al., 2019] Matuszewski, B., Schnell, N., and Bevilacqua, F. (2019). Interaction topologies in mobile-based situated networked music systems. *Wireless Communications and Mobile Computing*, 2019. (Cited on 16, 108.)

[McKinney, 2014] McKinney, C. (2014). Quick live coding collaboration in the web browser. In *NIME*, pages 379–382. (Cited on 16.)

[McLeod, 2001] McLeod, K. (2001). Genres, subgenres, sub-subgenres and more: Musical and social differentiation within electronic/dance music communities. *Journal of Popular Music Studies*, 13(1):59–75. (Cited on 125.)

[Miell et al., 2005] Miell, D., MacDonald, R. A., Hargreaves, D. J., Pavlicevic, M., Wilson, G., Powles, J., Tolchinsky, D. E., Glennie, E., et al. (2005). *Musical communication*. Oxford University Press on Demand. (Cited on 156.)

[Minsky, 1974] Minsky, M. (1974). A framework for representing knowledge. In *Computation & intelligence*. (Cited on 124.)

[Miranda and Wanderley, 2006] Miranda, E. R. and Wanderley, M. M. (2006). *New digital musical instruments: control and interaction beyond the keyboard*, volume 21. AR Editions, Inc. (Cited on 15.)

[Misdariis, 2018] Misdariis, N. (2018). Sciences du design sonore-approche intégrée du design sonore au sein de la recherche en design.(sciences of sound design-an integrated approach of sound design within the design research framework). (Cited on 156.)

[Mittelstadt et al., 2016] Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., and Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2):2053951716679679. (Cited on 149.)

[Mnih et al., 2015] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529. (Cited on 60, 64, 74.)

[Momeni and Henry, 2006] Momeni, A. and Henry, C. (2006). Dynamic independent mapping layers for concurrent control of audio and video synthesis. *Computer Music Journal*, 30(1):49–66. (Cited on 30.)

[Monache et al., 2010] Monache, S. D., Polotti, P., and Rocchesso, D. (2010). A toolkit for explorations in sonic interaction design. In *Proceedings of the 5th audio mostly conference: a conference on interaction with sound*, page 1. ACM. (Cited on 44, 61.)

[Monjou, 2014] Monjou, M. (2014). La recherche en design et sa publicité. In *Azimuts 40-41*. (Cited on 145.)

[Montecchio et al., 2019] Montecchio, N., Roy, P., and Pachet, F. (2019). The skipping behavior of users of music streaming services and its relation to musical structure. *arXiv preprint arXiv:1903.06008*. (Cited on 125.)

[Morris and Fiebrink, 2013] Morris, D. and Fiebrink, R. (2013). Using machine learning to support pedagogy in the arts. *Personal and ubiquitous computing*, 17(8):1631–1635. (Cited on 25, 45.)

[Murphy, 2012] Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press. (Cited on 4, 17, 127, 135.)

[Nika et al., 2017] Nika, J., Déguernel, K., Chemla, A., Vincent, E., Assayag, G., et al. (2017). Dyci2 agents: merging the" free"," reactive", and" scenario-based" music generation paradigms. In *International Computer Music Conference*. (Cited on 20, 45.)

[Niv, 2009] Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53(3):139–154. (Cited on 60.)

[Nymoen et al., 2011] Nymoen, K., Caramiaux, B., Kozak, M., and Torresen, J. (2011). Analyzing sound tracings: a multimodal approach to music information retrieval. In *Proceedings of the 1st international ACM workshop on Music information retrieval with user-centered and multimodal strategies*, pages 39–44. ACM. (Cited on 12.)

[Pachet, 2003] Pachet, F. (2003). The continuator: Musical interaction with style. *Journal of New Music Research*, 32(3):333–341. (Cited on 45.)

[Pachet et al., 2013] Pachet, F., Roy, P., Moreira, J., and d'Inverno, M. (2013). Reflexive loopers for solo musical improvisation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2205–2208. ACM. (Cited on 20, 87.)

[Paine, 2002] Paine, G. (2002). Interactivity, where to from here? *Organised Sound*, 7(3):295–304. (Cited on 127.)

[Palmer, 1997] Palmer, C. (1997). Music performance. *Annual review of psychology*, 48(1):115–138. (Cited on 147.)

[Pasquier et al., 2016] Pasquier, P., Eigenfeldt, A., Bown, O., and Dubnov, S. (2016). An introduction to musical metacreation. *Computers in Entertainment (CIE)*, 14(2):2. (Cited on 123.)

[Patel et al., 2011] Patel, K., Drucker, S. M., Fogarty, J., Kapoor, A., and Tan, D. S. (2011). Using multiple models to understand data. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, page 1723. (Cited on 63.)

[Petitmengin, 2006] Petitmengin, C. (2006). Describing one's subjective experience in the second person: An interview method for the science of consciousness. *Phenomenology and the Cognitive sciences*, 5(3-4):229–269. (Cited on 67.)

[Phillips-Silver et al., 2010] Phillips-Silver, J., Aktipis, C. A., and Bryant, G. A. (2010). The ecology of entrainment: Foundations of coordinated rhythmic movement. *Music Perception: An Interdisciplinary Journal*, 28(1):3–14. (Cited on 132.)

[Rajaonarivo et al., 2017] Rajaonarivo, L., Courgeon, M., Maisel, E., and De Loor, P. (2017). Inline co-evolution between users and information presentation for data exploration. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*, pages 215–219. ACM. (Cited on 64.)

[Ramos et al., 2019] Ramos, G., Suh, J., Ghorashi, S., Meek, C., Banks, R., Amershi, S., Fiebrink, R., Smith-Renner, A., and Bansal, G. (2019). Emerging perspectives in human-centered machine learning. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, page W11. ACM. (Cited on 2, 127.)

[Resnick, 2007] Resnick, M. (2007). All i really need to know (about creative thinking) i learned (by studying how children learn) in kindergarten. In *Proceedings of the 6th ACM SIGCHI conference on Creativity & cognition*, pages 1–6. ACM. (Cited on 3, 88.)

[Resnick et al., 2005] Resnick, M., Myers, B., Nakakoji, K., Shneiderman, B., Pausch, R., Selker, T., and Eisenberg, M. (2005). Design principles for tools to support creative thinking. *Working Paper*. (Cited on 22, 57, 58, 61, 62, 65, 87.)

[Ricci et al., 2011] Ricci, F., Rokach, L., and Shapira, B. (2011). Introduction to recommender systems handbook. In *Recommender systems handbook*, pages 1–35. Springer. (Cited on 19.)

[Richardson et al., 2019] Richardson, R., Schultz, J., and Crawford, K. (2019). Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *New York University Law Review Online, Forthcoming*. (Cited on 149.)

[Risset, 2004] Risset, J.-C. (2004). Évolution des outils de création sonore. *Champs culturels*. (Cited on 14.)

[Risset et al., 2007] Risset, J.-C. et al. (2007). Fifty years of digital sound for music. In *Proceedings of the 4th Sound and Music Computing Conference (SMC)*. (Cited on 96.)

[Risset and Wessel, 1999] Risset, J.-C. and Wessel, D. L. (1999). Exploration of timbre by analysis and synthesis. In *The psychology of music*, pages 113–169. Elsevier. (Cited on 105.)

[Rittel, 1972] Rittel, H. W. (1972). *On the Planning Crisis: Systems Analysis of the "First and Second Generations"*. Institute of Urban and Regional Development. (Cited on 63, 145.)

[Roberts et al., 2019] Roberts, A., Engel, J., Mann, Y., Gillick, J., Kayacik, C., Nørly, S., Dinculescu, M., Radebaugh, C., Hawthorne, C., and Eck, D. (2019). Magenta studio: Augmenting creativity with deep learning in ableton live. In *Proceedings of the International Workshop on Musical Metacreation (MUME)*. (Cited on 19.)

[Rodet et al., 1987] Rodet, X., Depalle, P., and Poirot, G. (1987). Speech analysis and synthesis methods based on spectral envelopes and voiced/unvoiced functions. In *European Conference on Speech Technology*. (Cited on 97.)

[Rowe, 1992] Rowe, R. (1992). *Interactive music systems: machine listening and composing*. MIT press. (Cited on 13.)

[Ruotsalo et al., 2014] Ruotsalo, T., Jacucci, G., Myllymäki, P., and Kaski, S. (2014). Interactive intent modeling: Information discovery beyond search. *Commun. ACM*, 58(1):86–92. (Cited on 64.)

[Ruthmann, 2007] Ruthmann, S. A. (2007). Strategies for supporting music learning through online collaborative technologies. *Music education with digital technology*, pages 131–141. (Cited on 16.)

[Sawyer, 2014] Sawyer, R. K. (2014). *Group creativity: Music, theater, collaboration*. Psychology Press. (Cited on 13.)

[Schacher et al., 2014] Schacher, J. C., Bisig, D., and Kocher, P. (2014). The map and the flock: Emergence in mapping with swarm algorithms. *Computer Music Journal*. (Cited on 30.)

[Schaeffer, 1966] Schaeffer, P. (1966). *Traité des objets musicaux*. Le Seuil. (Cited on 13.)

[Schlaug et al., 2005] Schlaug, G., Norton, A., Overy, K., Winner, E., et al. (2005). Effects of music training on the child's brain and cognitive development. *ANNALS-NEW YORK ACADEMY OF SCIENCES*, 1060:219. (Cited on 13.)

[Schnell et al., 2017] Schnell, N., Matuszewski, B., Lambert, J.-P., Robaszkiewicz, S., Mubarak, O., Cunin, D., Bianchini, S., Boissarie, X., and Cieslik, G. (2017). Collective loops: multimodal interactions through co-located mobile devices and synchronized audiovisual rendering based on web standards. In *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*, pages 217–224. ACM. (Cited on 109.)

[Schnell et al., 2009] Schnell, N., Röbel, A., Schwarz, D., Peeters, G., Borghesi, R., et al. (2009). *MuBu and friends–Assembling tools for content based real-time interactive audio processing in Max/MSP*. Ann Arbor, MI: Michigan Publishing, University of Michigan Library. (Cited on 33.)

[Schön, 2017] Schön, D. A. (2017). *The reflective practitioner: How professionals think in action*. Routledge. (Cited on 146.)

[Schuler and Namioka, 1993] Schuler, D. and Namioka, A. (1993). *Participatory design: Principles and practices*. CRC Press. (Cited on 24, 107.)

[Schwarz et al., 2006] Schwarz, D., Beller, G., Verbrugghe, B., and Britton, S. (2006). Real-time corpus-based concatenative synthesis with catart. In *9th International Conference on Digital Audio Effects (DAFx)*, pages 279–282. (Cited on 35, 43.)

[Schwarz and Schnell, 2009] Schwarz, D. and Schnell, N. (2009). Sound search by content-based navigation in large databases. In *Sound and Music Computing (SMC)*, pages 1–1. (Cited on 85.)

[Scurto, 2016] Scurto, H. (2016). Creative interaction paradigms for machine learning-based design of gestural musical interfaces,. Pre-doctoral report, Ecole Normale Supérieure de Cachan, Goldsmiths, University of London. (Cited on 2, 25, 150.)

[Scurto and Bevilacqua, 2017] Scurto, H. and Bevilacqua, F. (2017). Vimos: Enabling expressive mediation and generation of embodied musical interactions. Doctoral Symposium, 2017, 4th International Conference on Movement Computing (MOCO'17). (Cited on 11.)

[Scurto and Bevilacqua, 2018] Scurto, H. and Bevilacqua, F. (2018). Appropriating music computing practices through human-ai collaboration. In *Journées d'Informatique Musicale (JIM 2018)*. (Cited on 42.)

[Scurto et al., 2018a] Scurto, H., Bevilacqua, F., and Caramiaux, B. (2018a). Perceiving agent collaborative sonic exploration in interactive reinforcement learning. In *Proceedings of the 15th Sound and Music Computing Conference (SMC 2018)*. (Cited on 42, 86, 99.)

[Scurto et al., 2017] Scurto, H., Bevilacqua, F., and Françoise, J. (2017). Shaping and exploring interactive motion-sound mappings using online clustering techniques. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME'17)*. (Cited on 28, 45.)

[Scurto and Chemla, 2019] Scurto, H. and Chemla, A. (2019). Machine learning for computer music multidisciplinary research: A practical case study. In *Proc. of the 14th International Symposium on Computer Music Multidisciplinary Research, Marseille, France*. (Cited on 60, 96.)

[Scurto et al., 2016] Scurto, H., Fiebrink, R., et al. (2016). Grab-and-play mapping: Creative machine learning approaches for musical inclusion and exploration. In *Proceedings of the 2016 International Computer Music Conference*. (Cited on 25, 28, 29, 45, 63, 151.)

[Scurto et al., 2019a] Scurto, H., Liu, W., Matuszewski, B., Bevilacqua, F., Frechin, J.-L., Petrevski, U., and Schnell, N. (2019a). Entrain: Encouraging social interaction

in collective music making. In *ACM SIGGRAPH 2019 Studio*, page 5. ACM. (Cited on 108, 114.)

[Scurto et al., 2018b] Scurto, H., Van Kerrebroeck, B., and Bevilacqua, F. (2018b). Riding the co-explorers. In *19th International Society for Music Information Retrieval Conference (ISMIR 2018)*. (Cited on 60, 89.)

[Scurto et al., 2019b] Scurto, H., Van Kerrebroeck, B., Caramiaux, B., and Bevilacqua, F. (2019b). Designing deep reinforcement learning for human parameter exploration. *arXiv preprint arXiv:1907.00824*. (Cited on 60.)

[Sèdes et al., 2019] Sèdes, A., Assayag, G., and Desainte-Catherine, M. (2019). Informatique et musique : éloge de l'indisciplinarité. In *Symposium Les sciences de la musique: de nouveaux défis dans une société en mutation*. SFM-SFE-SFAM-AFIM. (Cited on 145.)

[Serra, 2017] Serra, X. (2017). The computational study of a musical culture through its digital traces. *Acta Musicologica*, 89(1):24–44. (Cited on 21.)

[Settles, 2010] Settles, B. (2010). Active learning literature survey. *University of Wisconsin, Madison*, 52(55-66):11. (Cited on 17, 88, 112.)

[Shahriari et al., 2016] Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. (2016). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175. (Cited on 86.)

[Shilman et al., 2006] Shilman, M., Tan, D. S., and Simard, P. (2006). Cuetip: a mixed-initiative interface for correcting handwriting errors. In *Proceedings of the 19th annual ACM symposium on User interface software and technology*, pages 323–332. (Cited on 63.)

[Shneiderman, 1999] Shneiderman, B. (1999). User interfaces for creativity support tools. In *Proceedings of the 3rd conference on Creativity & cognition*, pages 15–22. ACM. (Cited on 22.)

[Shneiderman, 2007] Shneiderman, B. (2007). Creativity support tools: Accelerating discovery and innovation. *Communications of the ACM*, 50(12):20–32. (Cited on 61, 62.)

[Silver et al., 2016] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484. (Cited on 60.)

[Simard et al., 2017] Simard, P. Y., Amershi, S., Chickering, D. M., Pelton, A. E., Ghorashi, S., Meek, C., Ramos, G., Suh, J., Verwey, J., Wang, M., et al. (2017). Machine teaching: A new paradigm for building machine learning systems. *arXiv preprint arXiv:1707.06742*. (Cited on 22.)

[Simondon, 1958] Simondon, G. (1958). *Du mode d'existence des objets techniques*. Éditions Aubier-Montaigne. (Cited on 124.)

[Sinclair, 2018] Sinclair, P. (2018). *Nouveaux auditoriums et pratiques d'écoute*. PhD thesis. (Cited on 146.)

[Small, 1998] Small, C. (1998). *Musicking: The meanings of performing and listening*. Wesleyan University Press. (Cited on 12, 132, 135, 156.)

[Spiegel, 1987] Spiegel, L. (1987). Regarding the historical public availability of intelligent instruments. *Computer Music Journal*, 11(3):7–9. (Cited on 15.)

[Spielmann, 2010] Spielmann, Y. (2010). *Video: the reflexive medium*. The MIT Press. (Cited on 20.)

[Steinbeis and Koelsch, 2007] Steinbeis, N. and Koelsch, S. (2007). Shared neural resources between music and language indicate semantic processing of musical tension-resolution patterns. *Cerebral Cortex*, 18(5):1169–1178. (Cited on 133.)

[Sterne, 2003] Sterne, J. (2003). *The audible past: Cultural origins of sound reproduction*. Duke University Press. (Cited on 13.)

[Stumpf et al., 2007] Stumpf, S., Rajaram, V., Li, L., Burnett, M., Dietterich, T., Sullivan, E., Drummond, R., and Herlocker, J. (2007). Toward harnessing user feedback for machine learning. In *Proceedings of the 12th international conference on Intelligent user interfaces*, pages 82–91. ACM. (Cited on 64.)

[Stumpf et al., 2009] Stumpf, S., Rajaram, V., Li, L., Wong, W.-K., Burnett, M., Dietterich, T., Sullivan, E., and Herlocker, J. (2009). Interacting meaningfully with machine learning systems: Three experiments. *International Journal of Human-Computer Studies*, 67(8):639–662. (Cited on 61, 64, 86.)

[Suresh and Guttag, 2019] Suresh, H. and Guttag, J. V. (2019). A framework for understanding unintended consequences of machine learning. *arXiv preprint arXiv:1901.10002*. (Cited on 88.)

[Surges and Dubnov, 2013] Surges, G. and Dubnov, S. (2013). Feature selection and composition using pyoracle. In *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*, page 19. (Cited on 136.)

[Sutton and Barto, 2011] Sutton, R. S. and Barto, A. G. (2011). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press. (Cited on 17, 46, 49, 57, 60, 61, 64, 66, 74, 75, 85, 86, 99, 135.)

[Tanaka and Donnarumma, 2019] Tanaka, A. and Donnarumma, M. (2019). The body as musical instrument. *The Oxford Handbook of Music and the Body*, page 79. (Cited on 101.)

[Taylor, 2014] Taylor, T. D. (2014). *Strange sounds: Music, technology and culture*. Routledge. (Cited on 13.)

[Thomaz and Breazeal, 2008] Thomaz, A. L. and Breazeal, C. (2008). Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172(6-7):716–737. (Cited on 22, 23, 46, 48, 60, 64, 80.)

[Thompson et al., 2019] Thompson, S., Scurto, H., and Fiebrink, R. (2019). Sound control: Supporting custom musical interface design for children with disabilities. In *Proceedings of the 19th International Conference on New Interfaces for Musical Expression (NIME 2019)*. (Cited on 25, 139, 152.)

[Torre et al., 2016] Torre, G., Andersen, K., and Baldé, F. (2016). The hands: The making of a digital musical instrument. *Computer Music Journal*, 40(2):22–34. (Cited on 14.)

[Van Kerrebroeck, 2018] Van Kerrebroeck, B. (2018). Deep reinforcement learning for the design of musical interaction. Master's thesis, Sorbonne Université. (Cited on 60, 74, 75, 99.)

[Vandenberghe and Slegers, 2016] Vandenberghe, B. and Slegers, K. (2016). Designing for others, and the trap of hci methods & practices. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 512–524. ACM. (Cited on 143.)

[Varni et al., 2011] Varni, G., Mancini, M., Volpe, G., and Camurri, A. (2011). A system for mobile active music listening based on social interaction and embodiment. *Mobile Networks and Applications*, 16(3):375–384. (Cited on 114.)

[Vial, 2015] Vial, S. (2015). Qu'est-ce que la recherche en design? introduction aux sciences du design. *Sciences du design*, 1:22–36. (Cited on 145.)

[Villani et al., 2018] Villani, C., Schoenauer, M., Bonnet, Y., Berthet, C., Cornut, A.-C., Levin, F., and Rondepierre, B. (2018). Donner un sens à l'intelligence artificielle. (Cited on 150.)

[Voillot et al., 2019] Voillot, M., Chevrier, J., Bevilacqua, F., and Eliot, C. (2019). Exploring embodied learning for early childhood education. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*, pages 747–750. ACM. (Cited on 26.)

[Wagstaff, 2012] Wagstaff, K. L. (2012). Machine learning that matters. In *In Procs of ICML*. Citeseer. (Cited on 21, 126.)

[Wanderley and Mackay, 2019] Wanderley, M. M. and Mackay, W. E. (2019). Hci, music and art: an interview with wendy mackay. In *New Directions in Music and Human-Computer Interaction*, pages 115–120. Springer. (Cited on 141.)

[Warnell et al., 2017] Warnell, G., Waytowich, N., Lawhern, V., and Stone, P. (2017). Deep tamer: Interactive agent shaping in high-dimensional state spaces. *arXiv preprint arXiv:1709.10163*. (Cited on 64, 74, 75, 86, 99.)

[Wiener, 1965] Wiener, N. (1965). *Cybernetics or Control and Communication in the Animal and the Machine*, volume 25. MIT press. (Cited on 105.)

[Wiener, 1988] Wiener, N. (1988). *The human use of human beings: Cybernetics and society*. Da Capo Press. (Cited on 21.)

[Wiggins, 2006] Wiggins, G. A. (2006). A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems*, 19(7):449–458. (Cited on 62.)

[Wikström, 2013] Wikström, P. (2013). *The music industry: Music in the cloud*. Polity. (Cited on 12, 19, 148.)

[Williams, 2016] Williams, A. (2016). Re-engineering hegemony: Glass bead in conversation with mat dryhurst, holly herndon and alex williams. *Glass Bead*, 1. (Cited on 136, 141.)

[Wishart, 1996] Wishart, T. (1996). *On sonic art*. Routledge. (Cited on 156.)

[Wong et al., 2011] Wong, W.-K., Oberst, I., Das, S., Moore, T., Stumpf, S., McIntosh, K., and Burnett, M. (2011). End-user feature labeling: A locally-weighted regression approach. In *Proceedings of the 16th international conference on Intelligent user interfaces*, pages 115–124. (Cited on 63.)

[Wright, 2005] Wright, M. (2005). Open sound control: an enabling technology for musical networking. *Organised Sound*, 10(3):193–200. (Cited on 15, 77.)

[Wright and Kanellopoulos, 2010] Wright, R. and Kanellopoulos, P. (2010). Informal music learning, improvisation and teacher education. *British Journal of Music Education*, 27(1):71–87. (Cited on 20.)

[Xenakis, 1992] Xenakis, I. (1992). *Formalized music: thought and mathematics in composition*. Pendragon Press. (Cited on 15.)

[Yang et al., 2018a] Yang, Q., Banovic, N., and Zimmerman, J. (2018a). Mapping machine learning advances from hci research to reveal starting places for design innovation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 130. ACM. (Cited on 62.)

[Yang et al., 2018b] Yang, Q., Scuito, A., Zimmerman, J., Forlizzi, J., and Steinfeld, A. (2018b). Investigating how experienced ux designers effectively work with machine learning. In *Proceedings of the 2018 Designing Interactive Systems Conference*, pages 585–596. ACM. (Cited on 62.)

[Yannakakis et al., 2014] Yannakakis, G. N., Liapis, A., and Alexopoulos, C. (2014). Mixed-initiative co-creativity. In *FDG*. (Cited on 63.)

[Ystad et al., 2019] Ystad, S., Aramaki, M., and Kronland-Martinet, R. (2019). Timbre from sound synthesis and high-level control perspectives. In *Timbre: Acoustics, Perception, and Cognition*, pages 361–389. Springer. (Cited on 98.)

[Zamborlin et al., 2014] Zamborlin, B., Bevilacqua, F., Gillies, M., and D'inverno, M. (2014). Fluid gesture interaction design: Applications of continuous recognition for the design of modern gestural interfaces. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 3(4):22. (Cited on 63.)

[Zappi and McPherson, 2014] Zappi, V. and McPherson, A. (2014). Dimensionality and appropriation in digital musical instrument design. In *NIME*, pages 455–460. (Cited on 43.)

[Zattra and Donin, 2016] Zattra, L. and Donin, N. (2016). A questionnaire-based investigation of the skills and roles of computer music designers. *Musicae Scientiae*, 20(3):436–456. (Cited on 156.)

[Zimmerman et al., 2007] Zimmerman, J., Forlizzi, J., and Evenson, S. (2007). Research through design as a method for interaction design research in hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 493–502. ACM. (Cited on 146.)

Hugo Scurto: *Designing With Machine Learning for Interactive Music Dispositifs*.
Final version as of January 2, 2020.

Supervisor: Frédéric Bevilacqua.