



**UNIVERSITY
OF ICELAND**

Hauntography

Algorithmic Presence and the Out of Joint in Music AI

Nicola Privato

Dissertation towards the degree of Doctor of Philosophy

2026

SCHOOL OF HUMANITIES

FACULTY OF ICELANDIC AND COMPARATIVE CULTURAL STUDIES

Íslensku og menningardeild Háskóla Íslands
hefur metið ritgerð þessa hæfa til varnar
við doktorspróf í menningarfræði

Reykjavík, 13. október 2026

Ásgrímur Angantýsson
Deildarforseti

Faculty of Icelandic and Comparative Cultural Studies
at the University of Iceland
has declared this dissertation eligible for a defence
leading to a Ph.D. degree in Cultural Studies

Doctoral Committee:
Þórhallur Magnússon, supervisor
Einar Torfi Einarsson, supervisor
Franziska Schroeder
Jón Ólafsson

Hauntography: Algorithmic Presence and the Out of Joint in Music AI

© Nicola Privato
Reykjavík 2026

Dissertation for a doctoral degree at the University of Iceland. All rights reserved.
No part of this publication may be reproduced in any form
without written permission of the author.

ISBN 978-9935-569-14-1
ORCID: 0009-0003-8316-312X

Abstract

Artificial Intelligence (AI) has developed at a rapid pace over the last decade. Accessible computing power and increasingly large datasets, coupled with recent advancements in machine learning techniques, have led to its extensive application in creative domains, including music and the arts. Among AI architectures, the deep learning subset has recently gained prominence. Deep learning models feature multiple processing layers that make them capable of identifying complex statistical patterns in large amounts of data. This architectural complexity makes these systems akin to black boxes, obscuring the processes that lead to the generation of their outputs. For this reason, explainable AI (XAI) has emerged as a prominent field of investigation. While XAI research often focuses on breaking open the AI black box, this dissertation takes a different stance and investigates the creative and epistemological value of AI's opacity, using it as a theoretical probe to query AI itself.

This dissertation asks (i) how the inherent opacity of deep learning systems can come to represent a creative resource, (ii) how this perspective can contribute to reassessing shifting notions of agency, authorship and intentionality in human-AI co-creation, and (iii) how insights from artistic practice may be used to understand the general impact of AI technologies on our society.

The dissertation introduces two new interfaces for musical expression (NIMEs) designed to perform with real-time neural audio synthesis (NAS), a deep learning synthesis technique that has recently come to the fore due to its high fidelity and low latency. Three original artistic works based on NAS are then offered and discussed. This corpus, together with works from other artists, is used to answer the first research question, illustrating how artists often leverage AI's opacity as a creative resource, and increasingly project presences of different kinds onto the hidden layers of AI models.

To mobilise the insights from practice and answer the second research question, the dissertation draws on hauntology theory, and through a practice-informed comparative reading of Jacques Derrida and Mark Fisher, develops a framework I term *hauntography*. Central to this framework are Fisher's notions of the *weird* and the *eerie*, which through Derrida I frame as two co-constitutive modes of phenomena viewed as spectral. Hauntography thus provides an extended vocabulary applicable to creative practice and critical analysis of phenomena from the perspective of epistemological and ontological liminality.

To answer the third research question, the theoretical framework is operationalised through thematic analysis of interviews with artists whose practices involve generative AI, and with whom I collaborated during the research. This analysis offers three interrelated insights. First, I observe how AI musical interfaces are often valued by artists insofar as they introduce a weird and eerie perspective on the sonic world represented by the training data. I then develop the notion of *spectrophonia* to describe how AI outputs bear the deferring traces of the source material, pulverised and reconfigured into something other, yet eerily familiar. Finally, I extend *spectrophonia* beyond music AI to reflect on AI more broadly. Through a hauntographic reading of Marx, supported by the interviews and by practice, I describe AI as a *spectral commodity*. While commodities typically conceal the social relations they reify, AI is structurally unable to do so, since these, inscribed in its training data, have to be re-enacted through use.

Útdráttur

Gervigreind hefur þróast á miklum hraða undanfarinn áratug. Aðgengilegt reikniafl og sífellt stærri gagnasöfn, ásamt stórstíguðum framförum á sviði vélanáms (e. *machine learning*), hafa leitt til víðtækrar notkunar gervigreindar í skapandi greinum, þar á meðal í tónlist og myndlist. Meðal gervigreindarkerfa hefur djúpnám (e. *deep learning*) nýlega rutt sér til rúms. Djúpnámslíkön einkennast af mörgum vinnslulögum sem gera þeim kleift að greina flókin tölfræðileg mynstur í miklu magni gagna. Þetta flækjustig gerir kerfin áþekkt svokölluðum „svörtum kössum“, þar sem ferlin sem leiða til framsetningar efnis eru hulin. Af þessum sökum hefur útskýranleg gervigreind (e. *Explainable AI, XAI*) nýlega orðið áberandi rannsóknarsvið. Þó að rannsóknir á XAI beinist oft að því að opna svarta kassann, tekur þessi ritgerð aðra stefnu. Í henni er sjónum beint að skapandi og þekkingarfræðilegu gildi ógagnsærrar gervigreindar, og notar í fræðilegri nálgun til að kanna gervigreindina sjálfa.

Þessi ritgerð spyr (i) hvernig innbyggt ógagnsæi djúpnámskerfa geti orðið uppspretta sköpunar, (ii) hvernig sú sýn geti stuðlað að endurmati á síbreytilegum hugmyndum um atbeina, höfundarvald og ásetning í samsköpun manns og gervigreindar, og (iii) hvernig innsýn úr listrænni iðkun geti nýst til að skilja almenn áhrif gervigreindartækni á samfélagið.

Ritgerðin gerir grein fyrir tveimur nýjum viðmótum til tónlistarsköpunar (e. *New Interfaces for Musical Expression, NIMES*) sem hönnuð eru til flutnings með tauganetstengdri hljóðgervingu í rauntíma (e. *Neural Audio Synthesis, NAS*), en sú djúpnámsaðferð hefur nýlega rutt sér til rúms vegna mikilla hljóðgæða og snerpu. Því næst eru sett fram og rædd þrjú frumsamin listaverk byggð á NAS. Þessi verk, ásamt verkum annarra listamanna, eru nýtt til að svara fyrstu rannsóknarspurningunni og sýna hvernig listamenn nýta oft ógagnsæi gervigreindar sem skapandi auðlind og varpa í auknum mæli ólíkum tegundum viðveru (e. *presence*) inn í hulin lög gervigreindarlíkana.

Til að virkja innsýnina úr listrænni iðkun og svara annarri rannsóknarspurningunni styðst ritgerðin við kenningar vofufræðinnar (e. *hauntology*) og þróar fræðilegan ramma sem ég nefni *vofuritun* (e. *hauntography*). Í kjarna þessa ramma eru hugtök Fishers um hið *undarlega* (e. *the weird*) og hið *óhugnanlega* (e. *the eerie*), sem ég set fram í gegnum kenningar Derrida sem tvö samverkandi birtingarform fyrirbæra er skilja má sem vofukennd (e. *spectral*). Vofuritun býður þannig upp á breikkaðan orðaforða fyrir skapandi iðkun og gagnrýna greiningu fyrirbæra út frá þekkingarfræðilegri og verufræðilegri jaðarstöðu.

Til að svara þriðju rannsóknarspurningunni er fræðilegi ramminn virkjaður í þemagreiningu á viðtölum við listamenn sem vinna með skapandi gervigreind og ég hef átt í samstarfi við á rannsóknartímabilinu. Þessi greining leiðir í ljós þrjár samtengdar niðurstöður. Í fyrsta lagi sýni ég hvernig listamenn meta gjarnan tónlistarviðmót byggð á gervigreind að því marki sem þau bjóða upp á undarlega og óhugnanlega sýn á þann hljóðheim sem þjálfunargögnin endurspegla. Í framhaldinu þróa ég hugtakið vofuhljómur (e. *spectrophonia*) til að lýsa því hvernig verk gervigreindarinnar bera með sér frestuð spor (e. *deferred traces*) hins upprunalega efnis, brotið niður og endurskipulagt í eitthvað annað, en þó óhugnanlega kunnuglegt. Að lokum beiti ég vofuhljómi út fyrir svið gervigreindartónlistar og ígrunda gervigreind í víðara samhengi. Með lestri á Marx að hætti vofuritunar, sem studdur er af viðtölum og listiðkun, geri ég grein fyrir gervigreind sem vofukenndum varningi. Rannsóknirnar sýna að þótt varningur hylji jafnan þau félagslegu tengsl sem hann hlutgerir, er gervigreindin ófær um það þar sem tengslin eru innbyggð í þjálfunargögnum hennar og því ávallt endurtekin í notkun.

Acknowledgements

This research project would not have been possible without the invaluable support from many colleagues, friends, and family members. Their insights, constructive critiques, and encouraging words have been crucial throughout this exciting and challenging journey.

I heartily thank my wonderful colleagues from the Intelligent Instruments Lab. I hear that PhDs are often solitary enterprises, and that this can be one of the most daunting aspects of doctoral work. This was not my experience. I was fortunate enough to share my days with a group of gentle people, off whom I could bounce new ideas, and engage in a reflexive practice that deeply enriched my critical and creative thinking.

My deepest thanks go to my colleagues and friends Victor Shepardson and Miguel Angel Crozzoli, with whom I shared years of exciting PhD research and from whose unique perspectives I learnt so much; to Halldór Úlfarsson, whose ability to turn all sorts of ideas into working prototypes astonishes me to this day; to Esther Þorvalds and Kamilla Henriau, whose help was crucial in settling in and finding my way through administrative requirements; to Stefanos Skialivas, whose warmth and kindness brightened the lab in this last year; to Jack Armitage and Sean Patrick O'Brien, who were always ready and excited to work on new challenges and ideas.

My thanks extend to the many postdoctoral researchers and visiting artists at the Intelligent Instruments Lab, whose perspectives were crucial in shaping the artistic and design outputs and the ideas expressed in this dissertation. I would like to thank in particular Giacomo Lepri, with whom I developed Stacco and shared the stage, and Robert Laidlow, Þórunn Björnsdóttir, Enrique Hurtado Mendieta, Erik Klinga, Voldemars Johansons, and Darija Andzakovic, who dedicated their time and attention to answering my questions.

A special thanks goes to my supervisors Thor Magnusson and Einar Torfi Einarsson, whose support, trust, and constructive critiques were invaluable. Leaving the comforts of one's home country to embark on such a life-changing adventure is challenging, but the support I received made the journey productive, exciting, and fun. I could have had no better supervision, and for this I am deeply grateful.

My thanks also extend to the entire doctoral committee, composed of Thor Magnusson, Einar Torfi Einarsson, Franziska Schroeder, and Jón Ólafsson, whose feedback and assistance contributed decisively to shaping this dissertation and the research from which it stems.

Finally, I would like to acknowledge my Italian family, especially Marco and Yoko, whom I could always count on in times of need, and my Costa Rican family. The peace of mind and space for thinking I found in the wonderful studio they graciously provided were invaluable to the completion of this work.

My most heartfelt gratitude goes to Natalia, my life partner, whose support made this journey imaginable, and to my son Mattia, who motivated me from the future-present. I look forward to our journey ahead.

Funding

This research is supported by the European Research Council (ERC) as part of the Intelligent Instruments project (INTENT), under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 101001848).

Declaration

I confirm that the research discussed in this dissertation is my own work. Where it has been carried out in collaboration with, or supported by, others, this is duly acknowledged. I attest that the work is original and does not infringe any third-party copyright or other intellectual property rights, or contain any confidential material. I acknowledge the use of Claude.ai (Anthropic) as an editing aid. The intellectual content and contributions are entirely my own. I confirm that this dissertation has not been previously submitted for the award of a degree by this or any other university.

Table of Contents

1	<i>Introduction</i>	2
1.1	Prologue	2
1.2	Research Overview	3
1.3	Research Questions.....	8
1.4	Chapters Overview.....	8
1.5	Positionality Statement.....	11
1.6	Dissemination	14
2	<i>Background</i>	16
2.1	Introduction	16
2.2	AI	17
2.2.1	Architectures and Models	18
2.2.2	Architectures, Models, Systems	21
2.3	AI Explainability	22
2.3.1	Interpretability, Explainability, and Beyond	22
2.3.2	Towards a Pragmatic View on XAI	24
2.3.3	XAIxArts.....	24
2.3.4	Interpreting Latent Space	25
2.3.5	Un-XAI	26
2.4	The Open Work.....	27
2.4.1	The Open Work and AI.....	28
2.5	Spectralities.....	30
2.5.1	Hauntology.....	31
2.5.2	Spectrality in Music.....	33
2.5.3	Sonic Hauntology	35
2.5.4	Spectrality and AI	36
2.5.5	A Genealogy of Hauntography	37
3	<i>Frameworks and Methods</i>	39
3.1	Introduction	39

3.2	Methods, Frameworks, Metaphors.....	40
3.2.1	Epistemological Metaphors	41
3.3	Practice Research.....	42
3.3.1	Subjectivity.....	43
3.3.2	Emergence	43
3.3.3	Circularity	44
3.3.4	Interdisciplinarity	45
3.4	Material Semiotics	45
3.4.1	Material Semiotics in Practice	47
3.4.2	A Hermeneutic Approach to Embodiment	48
3.5	Close Reading of Texts	50
3.5.1	Comparative Close Reading	50
3.5.2	Comparative Close Reading in Practice	51
3.5.3	Hauntographic Analysis of Collaborative Practice	53
3.6	Ethics.....	54
4	<i>Inherent Scores</i>	57
4.1	Introduction	57
4.2	Background	59
4.2.1	Inherent Scores.....	59
4.2.2	Composing with Real-time NAS	60
4.2.3	Magnetic Scores.....	61
4.3	Thales.....	63
4.3.1	Hands Holding the Void	64
4.3.2	Design	66
4.3.3	Sound Design	67
4.4	Stacco.....	68
4.4.1	Design	69
4.4.2	Sound Design	70
4.4.3	Embodied Sketching	71
4.5	Study 1 - Thales: XAI and Sound in Space	72
4.5.1	Research Ideation	73
4.5.2	DodekaOTTO.....	73
4.5.3	Insights.....	75

4.6	Study 2 - Stacco: XAI and Latent Regularisation.....	76
4.6.1	Study Design	77
4.6.2	Settings	78
4.6.3	Results.....	79
4.6.4	P2 and P4	80
4.6.5	Study Conclusions	81
4.7	From XAI to Un-XAI	82
4.7.1	Toward Spectrality	84
4.8	Conclusions	87
5	<i>Works in Movement</i>.....	89
5.1	Introduction	89
5.2	Mouja.....	91
5.2.1	Concept.....	92
5.2.2	Mouja.....	94
5.2.3	Performance Practice	97
5.2.4	Spectrality in Mouja.....	98
5.3	End of Empire	98
5.3.1	Context.....	100
5.3.2	AI and Empire.....	101
5.3.3	Methodology	102
5.3.4	Sound Design	103
5.3.5	Interaction	106
5.3.6	Spectrality in End of Empire	108
5.4	Mouja+.....	109
5.4.1	Text Scores and AI.....	111
5.4.2	System Design.....	112
5.4.3	Structure	114
5.4.4	Reception and Evaluation	116
5.4.5	Compositional Strategies.....	117
5.4.6	Dataset Exploration and Open Access.....	118
5.4.7	Spectrality in Mouja+	119
5.5	Haunted Work	120
5.5.1	Chunks, Feathers, Mirrors	121
5.5.2	The Open Work Gets Haunted.....	122

5.5.3	Co-creation, Performativity, Haunting	123
5.6	Haunting in Digital Culture	125
5.7	Conclusions	126
6	<i>Hauntography</i>.....	129
6.1	Introduction	129
6.2	A Hauntographic Glossary	131
6.3	Hauntology.....	134
6.3.1	Spectres at the End of History	134
6.3.2	The Mode of Being of a Spectre	135
6.3.3	Hauntology.....	138
6.3.4	Three Things of the Thing	140
6.4	Towards Sonic Hauntology	141
6.4.1	Capitalist Realism and Hauntological Resistance	143
6.4.2	Sonic Hauntology and the Metaphysics of Crackle	144
6.5	The Weird and the Eerie.....	145
6.5.1	The Weird	146
6.5.2	The Eerie	149
6.5.3	Convergences and Divergences.....	152
6.6	Hauntography.....	154
6.6.1	Heritage	155
6.6.2	Hauntographic Praxis.....	156
6.6.3	Creative Hauntography.....	157
6.6.4	Analytical Hauntography	158
6.7	Conclusions	159
7	<i>AI Hauntography</i>.....	161
7.1	Introduction	161
7.2	Hauntography Interviews.....	163
7.2.1	Hauntogram #1: Þóranna Björnsdóttir's Un-worldings.....	165
7.2.2	Hauntogram #2: Robert Laidlow and the Out Of	167
7.2.3	Hauntogram #3: Enrique Hurtado and the Cat	169
7.2.4	Hauntogram #4: Darija Andzakovic, the Brainchild and the Memory Palace	172
7.2.5	Hauntogram #5: Erik Klinga and Sonic Hauntology	174

7.2.6	Hauntogram #6: Voldemars Johansons and the In-between.....	175
7.2.7	Emerging Themes	177
7.3	The NIME at the End of the World	178
7.3.1	Weird Thresholds and Ground Truth.....	180
7.4	Spectrophonia	182
7.4.1	Defining Spectrophonia	182
7.4.2	Spectrophonia in Artistic Practice	184
7.5	(Another) Spectre of Marx.....	185
7.5.1	Commodities and Spectres	186
7.5.2	AI as a Spectral Commodity	188
7.5.3	Future Work.....	190
7.6	Conclusions	192
8	Conclusions	194
8.1	Primary Contributions.....	194
8.2	Secondary Contributions	196
8.2.1	AI NIMES.....	197
8.2.2	Artistic Works.....	197
8.2.3	Theoretical Contributions	198
8.3	Research Journey.....	199
8.4	Limitations.....	200
8.5	Future Work	201
8.6	Epilogue	203
9	References	204
10	Appendix	220
10.1	Publications	220
10.2	Performances	223
10.3	Installations and Exhibitions.....	225
10.4	Workshops	226
10.5	Open-source	227

List of Figures

Figure 1.1. Scramble GUI. Screenshot by author.....	12
Figure 4.1. Thales. Credits Esther Porvalds, 2023.....	63
Figure 4.2. Magnetic Score	64
Figure 4.3. From left to right: Giacometti’s Hands Holding the Void, Donnarumma’s Ominous, Thales	65
Figure 4.4. Thales’ magnetic interactions.....	67
Figure 4.5. Stacco V3.....	69
Figure 4.6. Embodied sketch on Stacco	72
Figure 4.7. Performing with Thales at Ars Electronica. Credits Bogi Nagy, 2023	74
Figure 4.8. Thales’ Icelandic spell score.....	86
Figure 5.1. <i>Mouja</i> at Fabryka Sztuki, Łódź. Credits Marta Zajac-Krysiak, 2024	94
Figure 5.2. Embodied sketch on Stacco	96
Figure 5.3. <i>End of Empire</i> . Eva Sajovic and Nicola Privato. Credits Shona Hamilton, 2023	100
Figure 5.4. <i>End of Empire</i> ’s technical details.	105
Figure 5.5. <i>End of Empire</i> columns 1–4.	106
Figure 5.6. Visitors interacting with <i>End of Empire</i>	108
Figure 5.7. <i>Mouja+</i> at SWRL Space. Credits Arturs Pavlovs, 2024.....	110
Figure 5.8. <i>Mouja+</i> GUI.....	114
Figure 5.9. <i>Mouja+</i> score display.....	115
Figure 5.10. Performing <i>Mouja+</i> . Credits Arturs Pavlovs, 2024	117
Figure 6.1. Club Silencio, <i>Mulholland Drive</i> . David Lynch, 2001.....	147
Figure 6.2. Bosque Quemado, Turrialba volcano, Costa Rica. Author’s photograph, 2025	150
Figure 6.3. <i>Useless Machine</i> . Retrieved from Wikipedia, Credits Mr. Donovan, 2020	158
Figure 7.1. Instruments developed at the Skerpla composition course	164
Figure 7.2. <i>Ashes</i> , Elfriede Jelinek, 2025.....	165
Figure 7.3. Recording of Robert Laidlow’s Techno-Utopia. Zubin Kanga on Stacco with BBC Orchestra. Credits Robin Clewley, 2025	168
Figure 7.4. Enrique Hurtado in Djúpavík. Credits Victor Shepardson, 2024	170
Figure 7.5. BYOV performing Sperger’s sonata at AIMC 2023. Credits Sean Patrick O’Brien, 2023	173
Figure 7.6. Voldemars Johansons at SWRL. Credits Arturs Pavlovs, 2024	176

List of Tables

Table 4.1. Study settings, rounds 1 and 2.....	79
---	----

There are ghosts in the machine,
and we are they, and they are we.

The Weird and the Eerie, Mark Fisher, 2016

1 Introduction

1.1 Prologue

Among the technological wonders presented at the 1939 New York World's Fair, the Voder, an invention by the Bell Labs engineer Homer Dudley, was received with particular curiosity and wonder (Dudley 1939). The device allows a human operator to produce realistic speech sounds, including vocal inflections and singing tones, through a system of keys and pedals replacing a vocoder's coding module. As Helen Harper, the Voder's operator in the 1939 demonstration film,¹ notes, mastering the instrument required over a year of dedicated practice.

I returned to the film of the Voder, which I remembered from my conservatory years, as I prepared a lecture on voice synthesis and electroacoustic music. Encountering it in a moment of interest around themes of presence, absence, and liminality, especially as applied to music technology, felt like experiencing it anew. As I listened to the speakers' voices and followed a slideshow of black-and-white pictures of the Voder's controllers reminiscent of Michel Waisvisz's *The Hands* (Torre et al. 2016), and of the young operator sitting at the console with the wrists slightly flexed above the keyboard and surrounded by a curious crowd, I inevitably paid attention to the coming and going of the characters, to how their auras lingered, overlapped, or receded as the theatrical presentation progressed.

The extract opens with a voiceover introducing three characters: Garrett, a male figure in charge of explaining the functioning of the Voder, Helen Harper, the operator, and the Voder itself. Garrett asks Harper to utter the sentence "she saw me" with different inflections. The Voder articulates a monotone sentence, then repeats it with affirmative and interrogative tones. Garrett continues asking Harper to greet everybody on behalf of the Voder, at different pitch heights. One can only imagine the effect on the 1939 audience as the conversation begins to shift from Garrett and Harper to Garrett and the Voder, which is now described as a "comparatively young man," whose voice's shifts in pitch betray, according to Garrett, some "moments of uncertainty."

¹ The original 1939 VODER demonstration film was produced by Bell Telephone Laboratories. A preservation copy is available through the Internet Archive (Dudley 1939).

As the theatrical duo between Garrett and the Voder (now pictured as a confused and depressed adolescent) unfolds, Harper, the skilled trainee and operator-ventriloquist, recedes inside the machine's depths. She resurfaces a few minutes later, after the Voder transitions into an old man nostalgically singing *Auld Lang Syne*, to answer a direct question from Garrett. And as Harper emerges, the Voder-elderly reverts to the Voder-machine.

Years later and in a different investigation, Derrida would describe the process whereby a machine "spiritises" itself while being an artifactual body, a puppet, and a sort of automaton, as the contradiction inherent to "automatic autonomy, mechanical freedom, technical life" (Derrida 1994a, 192). This "spiritisation" of the Voder is facilitated by a concealed and concealing accomplice, namely, the recording medium. The "medium of the media" (Derrida 1994a, 63) displaces the scene in time and space, introducing an informational bottleneck that frees the human and machinic characters from their physical constraints, turning them into spectres to the ears and eyes of the listeners: absences marked by language and by some kind of *work*. As Derrida puts it, by a certain power of transformation (Derrida 1994a, 9). This power (which, as will be argued, is the power of ghosts) endows the Voder and Helen Harper with the ability to emerge from and plunge into machinic depths without taking a single step, imprinting traces of their presence and absence on each other.

Harper the typist, the ventriloquist, the cog in the Voder's apparatus; the Voder-machine, the Voder-adolescent, the Voder-oldster; Garrett the chief, the comedian, the anchorman, who, as in Lynch's *Mulholland Drive* (2001), enunciates emphatically: "No hay banda!" Spectres keep multiplying in this theatre-of-three, not only in space but in time as well, whose edges they join into the shape of an ouroboros. Indeed, the Voder's identity crisis and the ghost of Helen Harper, invisibilised by her boss behind the theatrically spiritised machine, haunt the viewer now even more than they did then; and through the ears of our haunted present, even the Voder's rendition of *Auld Lang Syne* on the eve of WWII resonates as the omen of the impending conflict.

1.2 Research Overview

This dissertation, as the prologue anticipates, deals with spectres of a particular sort. Not with the paranormal dimensions of the otherworldly and phantasmatic, nor with well-established

traditions of electroacoustic music, but with the spectres that from a place of epistemological and ontological liminality haunt history, media, and technics (see Section 2.5).

As shown in the prologue, these spectres never rest. They leak through the joints of time, they haunt human and machinic bodies and talk through them, they appear when conjured away and return speech on their own terms. They perform their transformations from the rising of the curtain, from the depths of what cannot be fully known because unfathomable, because yet to come, because never fully *here*.

“One has to know. One has to know it. One has to have knowledge,” writes Derrida (Derrida 1994a, 9). Even as one plays an instrument, and especially as one does not fully understand its workings, at least at some level, there exists a certain urge, even an anxiety, about knowing. Not-knowing, in other words, acts as a vacuum. It wants to be filled with something or some-one. As Jonathan Sterne argues in discussing the fetishistic character of music technologies, when the inner workings of a musical instrument are not knowable, they turn into a vast and complex world, haunted by forces we do not comprehend, and whose powers they promise to grant access to (Sterne 2019).

Yet once these forces are evoked from unknowable depths, they reciprocate our gaze. Derrida describes this haunting reversal as the *visor effect* (Derrida 1994a, 125), the seeing from that which cannot be known, and the asymmetry of being seen without seeing. As I discuss in Chapter 5, artists working with AI technologies increasingly conjure a similar phenomenon through a tacit question, one that Jennifer Walshe puts in plain language as she urges: “who is listening, and how?” (Walshe 2023). This question neither calls nor cares for a literal answer, I believe. Rather, its intent is creative. To paraphrase Derrida, it marks the name, and thus the presence, of the listener (Derrida 1994a, 9). This effective transposition of the visor effect in music AI, let us term it here *auditor effect*, alludes to the relevance of engaging critically with the ghosts-in-the-machine, of returning speech to them especially when their language is unknown. As will be argued, not because “one has to know,” but because “we are they, and they are we” (Fisher 2016, 109).

Indeed, in his prescient reflections on generative technologies, Italo Calvino discusses how algorithmic recombination processes might allow us to step outside normative thinking and there encounter the hidden ghosts of the individuals and of their society (Calvino 1987, 20). As discussed in this dissertation, artists working with AI enjoy diving into these depths, conjuring presences they (we) playfully entangle with, they (we) question not to understand

the machine-in-itself, but rather to understand the machines-in-ourselves.² Scholarly work, I maintain, ought to follow this lead.

The creative potential of what returns our gaze from a place of liminality leads us to one knowledge gap that this dissertation addresses. While much research focuses on explaining AI from techno-centred or pragmatic perspectives, and typically positively, little scholarly work has been done in turning AI's inherent opacity, its unfathomability (fittingly, from the *fathom*, a measure of oceanic depths), into an epistemological resource rather than an AI explainability (XAI) conundrum. In this dissertation, I investigate this gap by asking *how AI's opacity enters the discourse, practices, and aesthetics in artistic creation and performance*.

The analysis of my own artistic and design practice, and of outputs from other artists engaging with music AI, shows how poetic meaning increasingly relies on the conjuring of liminal presences from the hidden layers of deep learning models. In other words, the deliberate openness that characterises contemporary art, and that Umberto Eco articulates with the term *open work* (Eco 1989), with generative AI brings along a question of presence. And as the open work gets increasingly haunted, it becomes crucial to develop frameworks capable of querying its ghosts.

I identify such a framework in hauntology theory, as framed by Jacques Derrida, whose seminal work *Specters of Marx* (1994a) is pivotal to the “spectral turn” characterising media studies, anthropology, sociology and cultural studies at the end of the twentieth century (Blanco and Peeren 2013). As this dissertation argues, hauntological thinking might contribute to addressing what Sturm et al. (2024) describe as a lack of adequate vocabularies for understanding AI's cultural operations beyond technical and functional framings, and the need for “more appropriate metaphors and concepts to critically analyse AI music, the existing music practices in which it takes part, and the new practices that might come from AI technology.”

From the first research objective thus derives a second one, namely understanding *how the discursive metaphors of spectrality and haunting, which emerged in addressing the first research question, might contribute to interrogating shifting notions of agency, authorship, and intentionality in human-AI co-creation*. I began addressing this question by applying hauntological thinking to the analysis of my own artistic practice. Yet, as the research progressed, I realised that the purposely ambiguous, sometimes indeterminate nature of

² I offer two interpretations of this sentence, a weak and a strong one. The weak interpretation highlights the entanglement between human and technical bodies, and the prosthetic nature of (music) technology. The strong interpretation points to hauntography's de-anthropomorphising stance, which views the human from the perspective of the non-human (see Section 6.5).

hauntology's notions constituted an operational obstacle in developing a hauntology-informed practice and expanding this framework into an instrument for analysis.

Mark Fisher's critique of Derrida epitomises this operational limitation. The thinker argues that Derrida's philosophical project, in which hauntology represents yet another deconstructive tool, "installed itself as a pious cult of indeterminacy, which at its worst made a lawyerly virtue of avoiding any definitive claim" (Fisher 2013a, 25). Yet, this critical stance did not prevent Fisher from grounding his analysis of cultural phenomena, and music in particular, upon the oblique notions introduced by Derrida, and from there attempting a crucial pragmatic move, through which hauntology becomes a tool for artistic intervention.

Fisher's contributions encompass an exceptionally insightful body of work engaging with music, mediation, spectrality, and the *strange*, which critically reflects on Western culture within neoliberalism (Fisher 2010; 2013a; 2013b; 2016). While Fisher's *sonic hauntology* (see Section 2.5.3) explicitly draws from Derrida, in reading Fisher's broad corpus one is nonetheless struck by the pull that Derrida's *Specters of Marx* seems to exert throughout. As I argue in Chapter 6, Fisherian themes such as *capitalist realism*, the *strange*, the *weird* and the *eerie*, acquire (if necessary) even more cohesiveness if interpreted through the lens of Derrida's seminal work.

Building on Derrida's hauntology and Fisher's operationalisation and further extension of it, in Chapter 6 I thus embark on the task of developing a vocabulary to delineate and analyse spectral phenomena without compromising their ambiguous ontological (thus, hauntological) status, from which stems the theoretical framework I term *hauntography*. This framework, which benefits science-technology studies and cultural studies in general, but AI musicology in particular, is developed through a historically situated comparative reading of the two thinkers' work.

Hauntography arguably represents the main contribution stemming from this dissertation, of which a complete description and vocabulary will be provided in Chapter 6. Nonetheless, as a way of establishing the intended outcome from the very beginning, I offer here a definition. Hauntography is *a posthumanist framework for the analysis and mobilisation of the traces (the spectres) of what operates by being absent (the ghost) and is apprehended through the affective and fictional categories of the weird and the eerie.*³

3 There exist multiple applications of the term *hauntography* in recent scholarly and fictional literature, each contributing a different definition and use. For an overview, see Section 2.5.5.

By extending Fisher's characterisation of the *strange* through a hauntological lens, hauntography identifies two co-constitutive modes (the *weird* and the *eerie*), which function as affective and fictional modalities. From this derives a twofold application of the framework: first, in shedding light on the phenomenological encounter with phenomena viewed as spectral; second, in offering a point of departure for a creative practice. These two hauntographic modalities, as Chapter 6 discusses, describe a circle of praxis that combines artistic research and critical reflection.

Along with the *weird* and the *eerie*, hauntography also introduces a series of notions and principles to guide the researcher. Among these are precise characterisations of *ghosts*, *spectres* and *spectrality*, the notions of *threshold*, *exorcism*, *hauntogram*, *auditor effect*, and a *principle of radical alterity* that informs the researcher's positionality from a non-anthropocentric perspective. These are discussed in depth in Chapter 6, and a glossary is provided in Section 6.2.

Hauntography sheds light on the entangled encounters between artists and AI systems but also speaks to the culture in which these encounters occur. The final contribution that this dissertation offers thus applies hauntography as an instrument of thematic and critical analysis within and beyond artistic practice. In the concluding chapter of the dissertation, I address the following question: *how do the findings in the intimate human-technology relationship of musical practice aid us in understanding the general impact of AI technologies on our society?* To answer this question, I offer a hauntography-informed deductive thematic analysis of interviews with artists I collaborated with over the past three years, whose works focus on diverse architectures, procedures, and artistic questions around music AI. The findings from the analysis prompt a concluding reflection on music AI, a reflection that is also sustained by my own practice research and by countless conversations with colleagues from the Intelligent Instruments Lab.

The answers to this third and last research question are inevitably partial. Not because of a lack of insights, but because of the wide research trajectory it opens, and that I hope to continue pursuing in the years to come. Nonetheless, the analysis confirms the value and actionability of the framework I offer by identifying a series of intertwined hauntographic themes, of which three, which I preview in Section 1.4, are expanded upon in Chapter 7.

1.3 Research Questions

The three research questions describe the dissertation's overarching trajectory encompassing the design of new interfaces for musical expression (NIMEs), artistic practice with music AI, comparative cultural studies and thematic analysis, and aim at understanding how AI enters and influences creativity and discourse.

- *RQ1: How does AI's opacity enter and influence the discourse, practices, and aesthetics in artistic creation and performance?*
- *RQ2: How can insights from these discursive practices contribute to developing tools to interrogate shifting notions of agency, authorship, and intentionality in human-AI co-creation?*
- *RQ3: How do findings in the intimate human-technology relationship of musical practice aid us in understanding the general impact of AI technologies on our society?*

1.4 Chapters Overview

The dissertation comprises eight chapters, of which Chapters 1 to 3 serve as introduction, background, and methodological framing. It then proceeds through four substantive chapters (Chapters 4 to 7), which may be viewed in continuity with each other, since together they describe the overarching narrative leading to the theorising and application of hauntography. More specifically, Chapters 4 and 5 address RQ1, Chapter 6 RQ2, and Chapter 7 RQ3.

Each chapter also introduces a series of secondary contributions that serve the main theoretical goal but also stand independently. In Chapter 8, I list both primary and secondary contributions, allowing readers to track the cumulative knowledge production across theoretical, artistic, and empirical domains.⁴

⁴ Because the interfaces and the performances represent secondary contributions, in the body of the dissertation I do not enter into the minute details of their design, different iterations, and source code. These are nevertheless available to the reader through the publication list and links provided in the dissertation.

Chapters 2 and 3: These introductory chapters outline the scholarly literature within which I situate the dissertation and describe the research methods applied in the investigation. The focus is on establishing a series of unifying theoretical and methodological underpinnings to the constellation of concepts and methods that the dissertation draws upon.

Chapter 2 overviews and examines literature on AI Explainability (XAI) outside and within the arts, and advances arguments for a pragmatic approach to XAI that accounts for the situated needs of the users in defining what, if anything, should be explained, and thus remains open to ambiguity. From here, the chapter discusses Umberto Eco's notion of the *open work* (Eco 1989), which reflects on openness in contemporary art, and, in a move that parallels the dissertation's narrative arc, expands openness and indeterminacy into the domains of presence, offering an overview of spectral theories in music, media studies, cultural studies, and other disciplines.

Chapter 3 positions hauntography, the main contribution of this dissertation, as a theoretical framework that can be applied in different fields through different methodologies. From here, the chapter organises the diverse research methods that underpin the interdisciplinary nature of this research around three clusters, namely practice research, material semiotics, and close reading of texts, discussing in detail how these are applied.

Chapter 4, Inherent Scores. Here, I introduce two novel NIMes I have developed during my research: Thales and Stacco. I describe them as *magnetic scores*, a particular type of *inherent score* (Tomás and Kaltenbrunner 2014) where the musical inscription emerges relationally through magnetic interactions. I designed these instruments to address the compositional challenges of neural audio synthesis (NAS) models (see Section 4.2.2). Through two empirical studies, one investigating spatial sound projection with Thales (Section 4.5), another examining latent space regularisation with Stacco (Section 4.6), I reflect on different ways of knowing a NAS model's latent space, and on how the unpredictability and complexity that XAI approaches often seek to regularise or disentangle are viewed by some as a creative dimension. The chapter concludes by using Derrida's notion of *exorc-analysis* (1994a, 58) to discuss how my AI interfaces, through magnetic actions-at-a-distance and seemingly autonomous behaviours, began staging a certain resistance to explanation.

Chapter 5, Works in Movement. This chapter introduces and analyses three artistic works. *Mouja* interprets real-time NAS through séance aesthetics. In this piece, I map the models' latent space with my AI interface Stacco, searching for impossible sounds of choirs with

infinite lungs and nonsensical voices. *End of Empire*, a large-scale, site-specific installation developed through participatory methods in Lancashire (UK), explores a non-extractivist use of NAS by engaging communities in data collection and sharing, a process that drastically shaped my understanding of the data. *Mouja+* introduces a pipeline for performing with text-to-audio models using scores from Fluxus artists, with sounds and texts conjuring onstage bodies and actions that never were.

I first discuss these three works individually, and continue reflecting on them through Eco's epistemological metaphor of the open work, which foregrounds the relevance of indeterminacy and openness in contemporary art. In *Mouja*, *End of Empire*, and *Mouja+*, though, openness and indeterminacy also foreground questions around presence. Following this insight, which I expand by looking at works of other artists engaging with music AI, I extend Eco's open work into what I term *haunted work*, a phrase identifying a growing corpus of artistic outputs in which algorithmic openness concretises around deferred presence.

Chapter 6, Hauntography. This chapter focuses on spectral theories, and in particular hauntology, to develop a framework around the findings presented in Chapters 4 and 5. Through a historically situated comparative reading of Derrida's *Specters of Marx* (1994a) and Fisher's body of work, particularly *The Weird and the Eerie* (2016), this chapter develops a vocabulary for analysing and mobilising phenomena, and AI in particular, through the lens of spectrality.

The chapter begins with Derrida's and Fisher's critique of the rhetoric of the *end of history*, of relevance to the Western philosophical discourse at the end of the twentieth century. Having established a common cultural and historical ground, I extend Derrida's notion of the *out of joint* through Fisher's notion of the *weird* and parallel the Derridean depictions of the ghost as *work* without presence with Fisher's notion of the *eerie*. The resulting framework, which includes a series of complementary notions developed throughout the chapter, provides a vocabulary for a praxis with the spectral in two modes: a *creative hauntography* for artistic production and an *analytical hauntography* for critical work.

Chapter 7, AI Hauntography. This chapter applies the hauntographic framework developed in Chapter 6 to music AI through deductive thematic analysis of interviews with artists whom my colleagues and I collaborated with at the Intelligent Instruments Lab. The analysis offers three insights.

First, I offer a view of AI NIMEs as *weird thresholds*, gates through which a *weird* perspective on the sonic world represented by the dataset infiltrates and destabilises it. Second, paralleling Schafer's *schizophonia* (Schafer 1969), I develop the notion of *spectrophonia* to describe how the dataset's multitude is pulverised and reconfigured by a deep learning model, producing sounds that carry echoes, or hauntographic traces, of their former selves. Third, I identify a countervailing dynamic in the reification of the model as agent, which obscures the social relations that spectrophonia brings to the foreground.

In the latter part of the chapter, I read the tension between spectrophonia and the model's reification through Marx's notion of commodity fetishism (Marx 1887/1995). Here, I argue that AI represents a uniquely *spectral commodity*, structurally unable to fully conceal the social relations it appropriates since it must re-enact them at each use. As an opening for future research, the chapter concludes by framing the current AI discourse as a form of cultural exorcism, an attempt to conjure away the revenants haunting the AI commodity.

1.5 Positionality Statement

My expertise lies primarily in music performance (idiomatic and non-idiomatic improvisation), jazz and electroacoustic composition (including design of interactive systems), and secondarily in the humanities. Before embarking on this project, I had already introduced deep learning tools of my design in my artistic practice, and published research on their artistic application. One of these systems, namely Scramble (Figure 1.1), combines Good Old-Fashioned AI (GOFAI) and deep learning techniques (see Section 2.2), allowing for the live training and dynamic combination of MIDI information either from files or from an instrument's live output (Privato et al. 2022).⁵

As my artistic practice with Scramble unfolded, it was inevitable to compare its affordances to those of other software and hardware instruments I perform with. Scramble was indeed highly unpredictable onstage, far more than any other tool I had experimented with. Oftentimes, it would converge on repetitive patterns around a piece's intervallic structures, only to abruptly break loose. On other occasions, it would wander through endless clusters of

⁵ Rather than using brute computational force as most current deep learning systems, Scramble encodes both low- and high-level musical relationships. It uses a series of Markov chains for pitch and an LSTM neural network for tempo, rhythm and dynamics, allowing musicians to combine, train, overlap in real time pitch models and rhythmic models from any MIDI file or sequence of files.

notes, evidently uninterested in generating anything meaningful to the ears of the listener. I never knew what to expect once on the stage, and whilst this unpredictability was challenging, it also came to represent the core matter of my artistic enjoyment.




Figure 1.1. Scramble GUI.

While inhabiting this space of uncertainty gradually became second nature for me, this was not always the case for the musicians I performed with: for each artist who enjoyed interacting with Scramble as much as I did, there was at least one who argued for a “taming” of my algorithmic companion’s wilderness. This divide led me to question whether and to what extent the unpredictability of my system, which like most deep learning systems is opaque to its designer almost as much as to those who engage with it, related to and could be modulated by addressing its opacity, and whether this would entail focusing on the system itself, on the epistemological and aesthetic assumptions of the musicians, or both. This reflection inevitably led to XAI territories and questions, and to this research project as initially formulated.

From the outset of my PhD research, I was aware of the importance of aligning my research questions with those of the Intelligent Instruments Lab, and more precisely with the second work package I joined. The aim of this work package is to develop a sociological theory of creative AI for human and non-human collaboration by reflecting through practice on the nature of creativity and distributed agency in intelligent systems, our musical relationships with non-human performers, and the public discourse surrounding these technologies. My education in modern cultures and literatures, which I pursued in parallel with the artistic one, has been crucial in this task, providing the necessary foundation for the theoretical excursions of Chapter 6.

Even more valuable has been the interweaving of practice and theoretical reflection that emerged through the combination of my artistic and humanistic backgrounds. Indeed, the pragmatic approach necessitated by practice, along with the engineering focus of colleagues within the NIME community, and the attentive oversight of my supervisors, helped prevent over-abstraction and ensured that insights were constantly evaluated empirically.

Finally, it is crucial to acknowledge the community in which my practice and investigation unfolded. The Intelligent Instruments Lab is home to an incredibly talented interdisciplinary group, and while this dissertation offers my personal contribution to knowledge, insights emerged through engaging in countless conversations, workshops, performances, and the collective sense-making processes that characterise daily life at the Lab.

One of the Intelligent Instruments Lab's priorities, toward which all members have actively worked, is applying citizen science methods (Riesch and Potter 2014) by inviting the Icelandic artistic community to respond to our research and sharing ideas and knowledge. Weekly, the Lab hosts open discussions with artists and practitioners interested in sharing their views. These events (more than one hundred and twenty over four years) profoundly shaped the ideas discussed in this dissertation.

The Intelligent Instruments Lab and the broader network of artists around it inevitably share certain beliefs, notions of validity, practices, and knowledge, and thus constitute what may be described as an epistemic community (Haas 1992). One might argue that this inevitably entails a certain collective bias. I acknowledge this, yet I also contend that knowledge is inevitably developed within epistemic communities. I do not embrace absolute epistemological relativism, but neither do I adhere to the foundationalist claim that knowledge can exist independently of the context that validates it. What matters, I believe, is acknowledging this situatedness and using it with care and attention to make sense of the reality we live in.

1.6 Dissemination

Various sections of this dissertation have been peer-reviewed and published in conference proceedings, journal articles, and book chapters.

Material from Chapters 2 and 3 has been published in the following conference paper and journal article, with the former expanded into a forthcoming book chapter on XAI in the arts:

- Nicola Privato and Jack Armitage (2023) “A Context-Sensitive Approach to XAI in Music Performance”, in *Proceedings of the First International Workshop on Explainable AI for the Arts, XAIxArts*.
- Nicola Privato and Thor Magnusson (2025) “The Open Work and AI: Co-creation and Haunting in Text Scores and Inherent Scores”, in *Organised Sound*. 2026:1-11. doi:10.1017/S135577182510085X

Parts of Chapter 4 have been published in:

- Nicola Privato, Thor Magnusson, and Einar Torfi Einarsson (2023) “Magnetic Interactions as a Somatosensory Interface”, in *Proceedings of the International Conference on New Interfaces for Musical Expression, NIME 2023*.
- Nicola Privato, Thor Magnusson and Einar Torfi Einarsson (2023) “The Magnetic Score: Somatosensory Inscriptions and Relational Design in the Instrument-Score”, in *Proceedings of the International Conference on Technologies for Music Notation and Representation, TENOR 2023*.
- Jack Armitage, Nicola Privato, Victor Shepardson, and Celeste Betancur Gutierrez (2023) “Explainable AI in Music Performance: Case Studies from Live Coding and Sound Spatialisation”, in *Proceedings of the Workshop “XAI in Action: Past, Present, and Future Applications”, part of Neural Information Processing Systems NeurIPS 2023*.
- Nicola Privato, Giacomo Lepri, Thor Magnusson and Einar Torfi Einarsson (2024) “Sketching Magnetic Interactions for Neural Synthesis”, in *Proceedings of the International Conference on Technologies for Music Notation and Representation, TENOR 2024*.

- Nicola Privato, Victor Shepardson, Giacomo Lepri and Thor Magnusson (2024) “Stacco: Exploring the Embodied Perception of Latent Representations in Neural Synthesis”, in *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME 2024.

Parts of Chapter 5 have been published in:

- Nicola Privato and Thor Magnusson (2025) “The Open Work and AI: Co-creation and Haunting in Text Scores and Inherent Scores”, in *Organised Sound*. 2026:1-11. doi:10.1017/S135577182510085X.
- Nicola Privato and Thor Magnusson (2025) “AI Reads Fluxus: Text Scores, Language Games, and the Art of Being Literal”, in *Proceedings of the International Conference on Technologies for Music Notation and Representation*, TENOR 2025.
- Nicola Privato (2024) “Mouja: Experiencing AI through Magnetic Interactions”, in *Proceedings of the 18th International Conference on Tangible, Embedded & Embodied Interaction*, TEI 2024.
- Giacomo Lepri and Nicola Privato (2024) “Magnetologues”, in *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME 2024.

Parts of Chapter 6 have been published in:

- Nicola Privato and Thor Magnusson (2024) “Querying the Ghost: AI Hauntography in NIME”, in *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME 2024.
- Nicola Privato and Thor Magnusson (2024) “AI Hauntology and the Hauntographic Method”, in *Proceedings of the 12th Conference on Computation, Communication, Aesthetics & X. Fabrica*, xCoAx 2024.

2 Background

2.1 Introduction

The dissertation's four substantive chapters (Chapters 4 to 7) draw on different fields of study and apply mixed research methodologies. This interdisciplinary effort is necessitated by the Intelligent Instruments Lab's research trajectory at the intersection of instrument design, artistic practice, and critical AI, and by the transversality of the research questions listed in the previous chapter. To accommodate this diversity, each chapter offers a background section contextualising locally relevant notions in depth. In this chapter, I provide the context for the overarching narrative that leads to hauntography's development and application, and critically position the dissertation in relation to it.

I begin by outlining the technical features of the AI models investigated in this dissertation, which I position within the generative subset of AI technologies. This technical foundation sets the stage for discussing the compositional and performative challenges addressed through practice, and that Section 4.2.2 details. These architectures do not represent the functioning of all AI systems, but offer concrete instances through which to understand broader patterns, in the same way as the broader discourse helps interpret the specific findings. In other words, generalisation is approached dialectically and with no universalising claims.

I then engage critically with the current discussions on AI Explainability (XAI), in particular within the arts, examining how XAI research tends to coalesce around two methodological clusters. Drawing on existing literature, I describe these as *techno-centred* and *pragmatic*. This distinction is necessary to introduce the first gap identified in Section 1.2, namely how XAI engages only limitedly with the creative and epistemic relevance of deep learning models' opacity, focusing instead on solving it *a priori*. This provides sufficient context for the practice research and qualitative studies in Chapter 4.

To ground the creative and epistemological relevance of what remains unexplained in AI, and thus open to interpretation, I turn to Eco's epistemological metaphor of the *open work*, describing the role of indeterminacy and openness in contemporary artistic production. This notion returns in Chapter 5, where it is argued that the space of openness and indeterminacy described by Eco, under the conditions introduced by generative technologies, often translates into artistic works reflecting on algorithmic presence. This consideration leads to the final

section of the chapter, in which I discuss spectral theories, provide a provisional definition of spectrality as discussed in this dissertation, and offer an overview of spectrality theory in multiple fields.

The aim of the overview on spectrality theory is not that of presenting all relevant literature around the theme, which is incredibly vast and transversal, and would thus require a dedicated monograph, but rather (i) of outlining the theoretical cluster that binds all chapters together, and (ii) of introducing the reader to the constellation of notions I draw upon throughout the dissertation. For instance, in the latter part of the chapter I introduce Fisher's and Derrida's framings of hauntology, which I draw upon in greater depth in Chapter 6; I also discuss R. Murray Schafer's notion of schizophonia (along with other *phonias*), which is revisited in Chapter 7 through a parallel with the hauntographic notion of *spectrophonia*.

However, this chapter does not introduce the entire theoretical corpus I engage with in the dissertation. The rationale for these exclusions is twofold: first, the locality of a certain framework or notion (e.g. Barad's agential realism, which is discussed tangentially in Chapter 4 in relation to my AI interface Thales; semiotics in HCI and NIME design, which is discussed in the methodology chapter; or the critical AI literature outlined in Chapter 5 to introduce one of my works); second, its dependency on this dissertation's contributions as foundational knowledge (e.g. Chapter 7's hauntographic framing of commodity fetishism, or the AI-based works of contemporary composers, which are drawn upon to extend insights from practice in Chapters 4 and 5).

2.2 AI

Over the past two decades, Machine Learning (ML) technologies have undergone significant advancements. The availability of large datasets, the scaling in computational resources, and the emergence of novel architectures and optimisation techniques have led artificial intelligence (AI) into a new phase of its development. This is particularly evident in the case of generative architectures, capable of learning probability distributions from large datasets to generate new material. These represent a subset of all AI systems, for which many taxonomies exist.

A first, broad taxonomical distinction may be drawn between GOFAI systems, which are typically rule-based and feature explicit symbolic representations, and generative

approaches that learn statistical patterns from more or less extended sets of data to generate new material (Boden 2014; Briot et al. 2017). Within generative music in particular, a crucial architectural classification distinguishes between neural audio synthesis (NAS) systems, which output continuous waveforms capturing the timbral qualities of the data they are trained upon, and symbolic music generation systems, which learn to generate discrete musical events represented by sequences of symbols (for instance, MIDI values) (Engel et al. 2017).

A different taxonomy divides AI systems by architectural depth. On the one hand are shallow systems, featuring no or a limited number of hidden layers between the input and output layers, such as simple multi-layer perceptrons (MLPs). These are typically trained on a limited amount of information, whether asynchronous or online. On the other hand are deep learning architectures, featuring a much higher number of layers, and typically performing hierarchical feature extraction from larger amounts of data (Schmidhuber 2014).

2.2.1 Architectures and Models

This dissertation focuses on a specific intersection of the above taxonomies: deep learning architectures trained on substantial datasets (whether relatively small, curated collections or large-scale archives, but in all cases much larger than those training shallow systems), that can be deployed in performative contexts for the generation of sound or complex musical events (therefore, not serving purely classification or regression tasks such as the Wekinator (Fiebrink et al. 2009), or FluCoMa (Tremblay et al. 2021)). More precisely, following Herremans et al.’s functional taxonomy, which organises systems according to the purposes for which they were designed, I focus on interactive systems employing an “online problem-solving approach, which can be real-time or not, to music generation that employs user input” (Herremans et al. 2018).⁶

I investigate three systems that afford different kinds of interactivity. I consider this diversity in modes of engagement particularly relevant for the purposes of this dissertation, in that the multiplicity of perspectives it grants provides more stable grounds to advance (careful and conditional) generalisations than a less diverse cluster of architectures, or the analysis of a single one. The three architectures I engage with are: RAVE, a real-time neural audio synthesis

⁶ In Chapter 7 the discussion expands from the machine learning systems detailed here to AI as a broader cultural phenomenon. Here as well, the focus will be on the recent AI wave of deep learning systems, rather than GOF AI ones.

system that incorporates continuous user input as parameters for latent space navigation and timbre transfer (Caillon and Esling 2021); Stable Audio Open, an asynchronous text-to-audio diffusion model that employs natural language user input to guide generation (Evans et al. 2024); and Notochord, a symbolic system designed for conversational co-improvisation that listens to a performer’s playing and responds in real time (Shepardson et al. 2022). These systems were chosen because of their diverse performative affordances (respectively, real-time latent manipulation, generation via textual prompts, conversational capabilities), operational temporalities, and training modalities.

RAVE: The Real-time Audio Variational Encoder (RAVE) is a neural audio synthesis method introduced by Caillon and Esling (2021). RAVE’s relatively high fidelity and low latency have drastically expanded the applications of neural synthesis in interactive contexts. The architecture learns by a two-phase procedure, consisting of a representation learning phase as a variational autoencoder (VAE) (Kingma and Welling 2019) followed by an adversarial fine-tuning phase which addresses sound quality. Training requires days of GPU computation (as a reference, on an NVIDIA A5000 training one model takes at least 78 hours), but after training, models can run in real time on laptop CPUs, for example through Pure Data, MaxMSP externals,⁷ SuperCollider UGens,⁸ and DAW VSTs,⁹ a feature that greatly facilitates RAVE’s integration within the setups of electroacoustic musicians.

As an autoencoder, RAVE consists of two main functions that the user can call separately: an encoding function that compresses a 48 kHz stream of audio to a stream of latent vectors in a multidimensional latent space, typically with a sampling rate of about 23 Hz and 4 to 32 latent dimensions, and a decoding function that synthesizes audio from latent vectors. A trained RAVE model is thus typically used in performance (i) by feeding a sound through both encoder and decoder to reinterpret it through the model’s training data, or (ii) by using control-rate signals to directly manipulate the latent dimensions before decoding. These two methods may also be combined, for instance, by using one’s voice as input and adding control signals to offset the encoder’s latent input values.

RAVE’s real-time capabilities and low latency, the relatively accessible training procedures, along with its ability to train on relatively small datasets, make it the ideal synthesis engine for NIME design. For instance, it is part of the synthesis engine of Shepardson’s Living

⁷ https://github.com/acids-ircam/nn_tilde

⁸ <https://github.com/elgiano/nn.ar>

⁹ <https://github.com/acids-ircam/rave>

Looper (Shepardson and Magnusson 2023), and its application is being investigated in the context of embedded synthesis (Pelinski et al. 2023). I chose this architecture for my AI interfaces Thales (Section 4.3) and Stacco (Section 4.4) and the installation *End of Empire* (Section 5.3). Similarly, this architecture has been chosen by many of the artists I collaborated with in the past three years, and whose interviews I analyse in Chapter 7. Besides artistic practice and NIME design, RAVE is also used in Chapter 4’s studies on AI explainability (Section 4.6).

Stable Audio: To expand beyond the specific affordances of RAVE, and particularly its dependence on relatively small datasets, I worked extensively with Stable Audio Open (Evans et al. 2024), a text-to-audio asynchronous diffusion model trained on a much larger dataset than RAVE, for the most part gathered from the open-source Freesound archive.¹⁰

Diffusion architectures such as Stable Audio Open generate sound files by progressively reversing a noise-adding process used for the training, and are typically conditioned on text vectors generated by LLMs. The user interacts by entering positive and negative prompts using natural language, and by defining parameters such as temperature and steps, corresponding to the randomness and the number of denoising iterations in the generation process.

Crucially, whereas real-time autoencoders such as RAVE continuously output small blocks of sound (typically, 2048 samples per block), thus allowing latent manipulation as the timbre transfer process unfolds, text-to-audio diffusion models generate one audio file per prompt and are thus asynchronous, their interactivity lying in how they interpret the user’s instruction into sound, rather than in their real-time response.

While limiting real-time manipulation, this asynchronous process enables macro formal control through natural language instructions. It may thus be said that RAVE and Stable Audio exhibit some complementary features, the former offering dynamic sound manipulation but no control over meso and macro formal aspects, and the latter capable of generating musical structures but offering no real-time control (Privato and Magnusson 2025). The performative pipeline I developed for *Mouja+*, and that I describe in Section 5.4, combines these two architectures into a single interactive system.

¹⁰ <https://freesound.org/>

Notochord: While most of the systems discussed in this dissertation synthesize raw sounds with no prior symbolic encoding of the musical information, Chapter 7 discusses insights from the use of Notochord, a symbolic conversational system developed by my colleague Victor Shepardson (Shepardson et al. 2022). By incorporating a symbolic model, I intended to expand the insights gathered from working with NAS systems beyond this domain.

Notochord is a deep probabilistic architecture designed for real-time MIDI performance with very low latency, featuring a recurrent neural network architecture with Gated Recurrent Units (GRU) layers. Its distinctive order-agnostic event organisation allows users to constrain any combination of sub-events (for instance, fixed pitch with the model choosing timing and instrument), thus enabling a set of diverse operations, for instance harmonisation or co-improvisation. Trained on the Lakh MIDI dataset (Raffel 2016),¹¹ Notochord handles polyphonic, multi-track MIDI generation while representing time in absolute seconds rather than tempo-relative units, and thus prioritises real-time responsiveness and interpretable interventions to explore embodied experiences of creative AI.

RAVE, Stable Audio Open, and Notochord span three different approaches to the performative deployment of generative AI in interactive contexts, including gestural sound manipulation, macro formal structuring through natural language, and conversational musical dialogue. Each system affords a distinct relationship between artist, interface, and generative architecture, thus offering a heterogeneous perspective on AI-mediated creation.

2.2.2 Architectures, Models, Systems

I conclude this section with a due terminological clarification. An aspect of the discourse concerning AI that I find particularly striking, and that I began to make sense of only retrospectively (indeed, we will return to this in Section 7.5), is that deep learning algorithms are typically described with the term *model* both before and after training. This evidently overlooks the fact that what allows them to meaningfully distribute the weights in their neurons, and thus enables the modelling, is the data they train upon.

To avoid the pitfall of conflating the ontologies of deep learning models prior to and after training, I will thus consistently use the terms *architecture* to refer to AI software before or independent of training, *model* to refer to the trained architecture, and *system* to refer to how

¹¹ <https://colinraffel.com/projects/lmd/>

AI is usually experienced in the world, that is, through some sort of interface, musical or otherwise, and in a certain context of use.

Having established the technical underpinnings that inform the dissertation, in the next sections I outline the theoretical background that informs this research. As the project progressed, the questions and theoretical framework have changed. Paralleling the dissertation’s structure, I thus account for this transition by accompanying the reader from Explainable AI (XAI) to a discussion on openness and indeterminacy in the arts, and finally into spectrality theory.

2.3 AI Explainability

The introduction of deep learning architectures notoriously poses a series of explainability challenges. This is primarily due to the complexity of their workings, which prevents inferences on the processes leading to the generation of their outputs (Linardatos et al. 2020). While the umbrella term for the burgeoning field of research that focuses on this problem is *Explainable AI (XAI)*, the literature often discusses the black box problem arising from these systems both as an explainability and an interpretability issue. These terms are often used interchangeably (Carvalho et al. 2019), and, as some argue, are often ill-defined (Lipton 2016).

2.3.1 Interpretability, Explainability, and Beyond

In the context of Multimodal Large Language Models (MLLMs), Dang et al. (2024) describe interpretable models as those possessing inherently understandable internal structures, and allowing for a straightforward comprehension of how inputs are transformed into outputs. Explainability, by contrast, involves post-hoc techniques that provide external analyses of a model’s decisions. Relatedly, Rudin et al. (2021) argue that interpretability focuses on designing models that are inherently understandable, whereas explainability aims at providing post-hoc explanations for existing black box models (for instance, models whose workings are incomprehensible to humans, or proprietary ones).

Though not unanimously accepted and applied, these characterisations enable advancing a twofold working distinction between the problems addressed by AI explainability

and interpretability research, which helps contextualise the nature of this contribution within the current XAI discourse. First, whereas explainability typically seeks to establish causal relations underlying a model’s behaviour, interpretability often focuses on designing systems whose logic is understood by the user at the outset. Second, whereas explainability entails post-hoc rationalisations, interpretability typically applies to the model in use.

Consider a VAE architecture such as RAVE. One approach to the explainability of this architecture would entail, for instance, mapping the hidden layers’ neural activations during performance to reveal post-hoc how and why certain sound features cluster through shared activation pathways. An interpretability approach, by contrast, would address the challenge of facilitating inference over the model’s latent distribution during interaction. This might involve adapting the latent space of a new model to align with one we are already familiar with (see Section 4.6), or organising the latent space around discrete sound descriptors (e.g., pitch, amplitude) as in the case study discussed in Section 2.3.4.

The terms *interpretability* and *explainability* typically refer to rational accounts of how a model operates. Yet, music practice often entails a direct and embodied interaction with a musical instrument or interface, with the reasons behind a certain organisation of sound features in an instrument imbued with AI being secondary to its potential for expression, discoverability, mystery, and magic; what Magnusson describes with the term *ergodynamics* (Magnusson 2019). In other words, musically relevant knowledge does not always need to be articulated, and is often tacit and embodied.

The research I present in Chapter 4 may be positioned within this relevant, and, in comparison to the XAI angles discussed in this section, underdeveloped perspective.¹² Crucially, to address XAI questions from this angle means shifting one’s view from top-down, a priori approaches to what constitutes a good explanation, and diving into the situated dialectic of accommodation and resistance that Pickering describes as the “mangle of practice” (Pickering 1995). As discussed in the next sections, the problem of addressing XAI questions pragmatically and situatedly goes to the heart of the current scholarly discussion, and especially of XAI application in the arts.

¹² Few scholars have investigated AI explainability through embodiment. Among these are Scurto and Postel’s latent sound walks (Scurto and Postel 2023), exploring a literal equivalence between RAVE’s latent space and 3D space as mediated by a virtual avatar, and the Intelligent Instruments Lab’s research on XAI and spatialisation (Armitage et al. 2025) and on the embodied perception of latent space discussed in Chapter 4 (Privato et al. 2024).

2.3.2 Towards a Pragmatic View on XAI

According to Lipton, explainability should not be viewed as a property that is specific to neural networks, as “sufficiently high dimensional models, unwieldy rule lists, and deep decision trees could all be considered less transparent than comparatively compact neural networks” (Lipton 2016, 5). Nor should it be viewed as a property of the architecture per se, since relatively large systems can be opaque to certain users and not others, depending on the background knowledge required to understand them (Páez 2020).

The phrase “I do not understand AI,” for instance, carries different meanings depending on who utters it, whether an AI engineer asking for technical clarifications, a sceptical customer who does not want to buy an intelligent fridge, or a scholar making a sarcastic remark at a conference on critical AI. Different types of explanation correspond to different *explanans*, for instance, facts or entities in the world (e.g., the cold explains the sore throat), theories or representations (e.g., infections are explained by some medical theory), or agents’ accounts (e.g., the doctor explaining how the vaccine works) (Potochnik 2016).

Accounting for this diversity of audiences and explanation types is crucial. As Miller puts it, “the very experts who understand decision-making models the best are not in the right position to judge the usefulness of explanations to lay users” (Miller 2019, 2). Yet, XAI research often adopts a top-down approach, using the intuition of the researcher rather than investigating what users value situationally, an aspect that also reflects a lack of literature around user-based measures, which limits the identification of correct problem formulations (Lipton 2016).

2.3.3 XAIxArts

In light of the critique described in Section 2.3.2, when investigating XAI questions in the arts it is important to reflect on what problems one aims to solve, and whether these respond to real needs. The diverse contributions from the 2023 workshop on XAI for the Arts investigate these and other questions, reframing the explainability problem from a much broader perspective, one that moves away from techno-centred, top-down assumptions on what is relevant to explain, and places the artist and the audience at the centre of the investigation (Bryan-Kinns et al. 2023b; Privato and Armitage 2023).

This reframing introduces several dimensions largely absent from technical XAI discourse. First, temporal considerations become particularly relevant. For instance, in a live performance real-time interpretation is crucial while post-hoc explanations are relevant as analytical tools. Second, explanations involve different stakeholders, such as composers, musicians, audiences, whose situated needs should be understood and responded to. Third, artists might value the model’s glitches and failures over the intelligibility of its workings. Indeed, as Arora and Sarkar argue, explanations are often ornamental to the artistic intention, and in the arts should be replaced with a notion of *sense-making*, understood as a relational and immanent “system of echoes, of resumptions and resonances” (Arora and Sarkar 2023).

On such theoretical foundations is grounded the work my colleague Jack Armitage and I contributed to on the occasion of the first XAIxArts workshop, which extends the pragmatic perspective advanced by Nyrup and Robinson in the context of medical AI (Nyrup and Robinson 2022) into music AI (Privato and Armitage 2023). Here, explanations are configured as speech acts (Austin 1975) that need to be evaluated in their ability to enable the audience to draw inferences, and in how well they satisfy the stakeholders’ subjective epistemic longing (Besold and Uckelman 2018, as cited in (Nyrup and Robinson 2022)).

In light of these reflections, the next section focuses on one particular XAI problem that my initial research drew upon, providing examples of how it has been previously addressed from a techno-centred angle, and advancing the rationale for its pragmatic investigation through artistic practice in Chapter 4.

2.3.4 *Interpreting Latent Space*

VAEs distribute sound features arbitrarily within the latent space, a process that Chapter 4 discusses and characterises as *arbitrary encoding* (see Section 4.2.2). This arbitrary organisation of sonic information hinders unambiguous mappings of latent dimensions to musical attributes.

This limitation has been addressed in different ways. One approach is to align latents with known audio descriptors, for instance by isolating certain features of the sound (e.g. pitch or amplitude) to be controlled explicitly, or by factoring the latent space into independent parts dealing with distinct aspects of sound (Devis et al. 2023). An example of the first approach in music AI may be found in Bryan-Kinns et al.’s work on explaining latent space, where the

authors map latent dimensions to meaningful musical parameters in a VAE model and provide feedback on distribution and position through a visual interface (Bryan-Kinns et al. 2023a).

While this method might enhance explainability in certain contexts, it nonetheless offers a top-down solution, with researchers defining what a good explanation is rather than the users. This echoes Morrison and McPherson’s diffractive dialogue, in which the two reflect on how MIDI technology provided a claviocentric ideological reservoir at a time when digital music hardware and software were attempting to move in new directions (Morrison and McPherson 2024). Translating this to music AI, in approaches attempting to discipline the entangled nature of VAEs’ latent space to a parametrised logic, it is easy to discern a certain *parametercentrism*, whereby conceptual structures and modes of interaction of modular synthesizers are applied to systems attempting once again to move in a new direction.¹³

Yet, as I show in Chapter 4, artists do not necessarily enjoy a highly interpretable latent space, and often value VAEs precisely due to the entangled and autonomous organisation of sound information they perform (see Section 7.2 for relevant examples). The study discussed in Section 4.6 operationalises discussions on how XAI questions and methods should account for the users’ needs and subjective epistemic longing, rather than imposing techno-centred assumptions on how users should or should not interact with VAEs.

2.3.5 *Un-XAI*

Through the pragmatic and embodied approach to XAI that Chapter 4 describes, I came to realise that under certain conditions artists consider a less interpretable latent space more valuable than a highly interpretable one. This insight was crucial to the research’s pivot into the meaning-making capacity of what remains open, indeterminate, of what resists explanation, as per the first research question (see Section 1.3).

The term *Un-XAI*, which titles this section, is not intended as a strict opposite to XAI. Rather than for theoretical scaffolding, I use it quite sparingly in the dissertation as an admittedly provocative acronym intended to remind myself, and perhaps fellow scholars invested in XAI in the arts, not to overlook the sense-making capacity of what, deliberately or not, resists explanation.

¹³ Through the framework I introduce in Chapter 6, this would be viewed as an act of *exorc-analysis*, whereby uncertainty is exorcised away by imposing pre-existing categories.

Another term that I enjoy using throughout the dissertation in relation to AI's opacity is *unfathomed*, which derives from the *fathom* as a unit for measuring the depth of water. This term projects not-knowing into a dimension of oceanic depth, echoing the extensiveness of the datasets on which AI architectures are trained, the inaccessibility of their hidden layers, and the presences that we project onto them.

To investigate this dimension of algorithmic unfathomability productively requires expanding beyond XAI's technical vocabulary into art aesthetics. The question of how indeterminacy generates meaning is indeed one that contemporary art has long addressed, and is at the core of Umberto Eco's seminal work *Opera Aperta* (Eco 1989). If looking at the unfathomed in AI implies dipping one's feet into an open ocean of emergent meanings, then Eco's open work provides a framework to dive into it. The following section contextualises and introduces Eco's theory of openness, which I apply in Chapters 4 and 5 to NIME designs and artistic work involving AI.

2.4 The Open Work

The second half of the twentieth century witnessed a significant transformation in compositional practices, as notation expanded beyond its traditional prescriptive function to explicitly incorporate openness and indeterminacy. This shift mirrored a broader evolution in the arts: as artists increasingly embraced relationality and emergence, the conventional notion of the work as a self-contained object representing the direct expression of the artist's inner intuition underwent a crisis. In the wake of this shift, the artwork became an increasingly open, indeterminate and interconnected system.

This transformation reflects a much broader, overdetermined cultural one. The rise of new media technologies, information theory and cybernetics fostered systemic thinking (Wiener 1961); postmodernism's rupture with tradition and scepticism toward grand narratives (Lyotard 1984), alongside theories prioritising the social construction of knowledge and the intersubjective nature of the aesthetic experience (Bourriaud 2009), contributed to the destabilisation of hierarchical structures and traditional modes of artistic enjoyment. These converging factors encouraged fluidity in interpretation, endowed audiences with an active role in shaping the work, and in music contributed to redefining the traditional roles of composers and performers.

To articulate this shift in opposition to Croce's aesthetics, which treats intuition and expression as indistinguishable, Umberto Eco introduced the notion of the *open work* (Eco 1989), an epistemological metaphor reflecting the new ways in which culture informs aesthetic interpretation and enjoyment. In Eco's theory, two complementary degrees of openness emerge: *openness of the first degree*, a cognitive process by which any work, including formally closed ones, generates new meanings through a continuous germination of internal relationships, and *openness of the second degree*, a deliberate poetics wherein indeterminacy and emergence are programmed within the very structure of the work.

This latter dimension has been central to musical experimentation in the second half of the twentieth century, with scores becoming "junctions of composer and performer creativity, or locations that [enable] the emergence of new and unforeseen musical results" (Magnusson 2019, 99). Artists realise this openness through diverse strategies: some focus on the emergent dimension of the inscription (Oliveros 2013), others problematise the score's situatedness in relation to the instrument (Tomás and Kaltenbrunner 2014), blur the role of composer, performer, and audience (Einarsson 2017), investigate live coding as a new path in the generation of the score (Magnusson 2011), or through data mining techniques incorporate social forces into the creative process (Privato and Novello 2021).

In Eco's framing of the open work, situatedness, interaction and entropy play a crucial role. The programmed interactivity requiring the listener to do some of the organising and structuring of the musical discourse defines the formally open, constantly evolving works that Eco describes as *works in movement*. Examples in this category are Cage's *4'33''*, where the work incorporates the sonic environment in which it is enacted, or Calder's *mobiles*, whose form is constantly reconfigured by shifting air currents. To discuss entropy, Eco draws on information theory, describing it as the degree of informational disorder inherent to the work, and arguing that "the richest form of communication (richest because most open) requires a delicate balance, permitting the merest order within the maximum disorder" (Eco 1989, 99). This view entrusts the artist with the role of defining, through a careful and deliberate use of structural openness and intention, the space within which poetic meaning arises.

2.4.1 *The Open Work and AI*

Eco primarily explores human and environmental factors as sources of openness, yet his drawing from information theory and cybernetics proved prescient. Since the writing of *Opera*

Aperta in 1962, where the metaphor of the open work was first introduced, generative systems have indeed evolved beyond mere recombination machines such as those imagined by Calvino (Calvino 1987) or employed by Xenakis in his stochastic works (Luque 2009), into machinic agents that artists at times describe as co-creative partners.¹⁴ Within this technical and aesthetic evolution, the open work's focus on situatedness and relationality, interaction, and entropy provides a fertile ground to reflect on the role of generative technologies in the arts.

For instance, in *Bias* (2021), a piece for clarinet and interactive music system, Artemi-Maria Gioti uses neural networks to simulate the composer's aesthetic judgement. Here, the model responds to the musician's live performance by imitation, by proposing new material, or by simply remaining silent. The composer establishes the system's architecture and sets in motion certain interaction rules, allowing the model's autonomous responses to progressively reveal biases, understood as the model's learnt patterns and, reflexively, the composer's own aesthetic assumptions. Through Eco's openness of the second degree, we can view the model as a co-creative partner that contributes to the performance while operating within constraints established by the artist; through the notion of openness of the first degree, we can observe the work as part of a network of relations in which aesthetic appreciation is relationally constructed, inviting, as David Robey puts it, "a collaboration and personal involvement that was never required by the traditional art of the past" (Eco 1989, xi).

This framework offers the advantage of maintaining working ontological distinctions, echoing Born and Barry's argument on the relevance of accounting for different scales of social relations and formations, as these mediate and are mediated by music (Born and Barry 2018). However, the open work functions primarily as a theory of aesthetics, even though one informed by political commitment, rather than as a theory of social mediation. Its relevance for this dissertation rather lies in the vocabulary of openness and indeterminacy it provides, and in the analysis of artistic poetics it is grounded upon, which together offer the theoretical and methodological point of departure to interrogate through artistic practice the new forms of indeterminacy introduced by AI.

For instance, Chapter 4 describes two AI NIMEs that build on the concept of inherent scores (where the notation is viewed as inherent to the circuitry of the interface (Tomás and Kaltenbrunner 2014)) to incorporate notational openness and reflect upon the compositional affordances introduced by real-time NAS systems. Similarly, Chapter 5 builds on Eco's

¹⁴ See Section 7.2 for examples of such descriptions.

framework to introduce three artistic works involving NAS, which I frame as *works in movement* in that their physical or algorithmic configurations are highly dependent on contextual re-arrangements.

What the open work does not offer, though, is the formalisation of a dimension of indeterminacy, that, while not specific to AI systems, appears to be increasingly relevant to artists working with generative AI. The argument I advance in Chapter 5 is that music AI's expansion of indeterminacy into the territories of presence requires a similar expansion of the epistemological metaphor advanced by Eco. I thus search for a lexicon in spectral theories, a body of critical work that understands culture, history, memory, technics through metaphors related to haunting and ghosts.

2.5 Spectralities

In the past three years, I engaged with spectrality theory extensively, through the available literature but also in my practice, and through countless discussions with fellow researchers and artists. It soon became apparent that the ambiguous ontological status of spectral phenomena is accompanied by a certain semantic uncertainty. To put it plainly, there is often no explicit consensus, especially in interdisciplinary fields, on what we mean by terms such as *ghosts*, *spectres*, or even by the notion of spectrality itself. Indeed, as Blanco and Peeren acknowledge in the title of their *Spectralities Reader* (and as I acknowledge in the title of this section), there exists no single spectrality (as a discipline or even as a field), because even when used as conceptual metaphors, ghosts remain “figure[s] of unruliness pointing to the tangibly ambiguous” (Blanco and Peeren 2013, 9).

While Chapter 6 fulfils the necessary task of establishing a framework to locally anchor the discussion and analysis, it is nonetheless possible and even crucial to establish two conceptual guard rails from the beginning. Firstly, as anticipated in Section 1.2, the spectrality I discuss in this dissertation does not deal with supernatural phenomena, ghostly apparitions and paranormal events per se, but rather includes them only insofar as these are viewed through the perspective of discourse. This positioning reflects a shift that occurred at the end of the twentieth century, when the figure of the ghost transitioned from an actual liminal entity or narrative artifice to an influential conceptual metaphor in global culture and academia (Blanco and Peeren 2013, 1).

Secondly, while there cannot exist one single definition of spectrality, more to theoretically anchor the dissertation's structure than to offer an impossibly exhaustive definition, I frame spectrality in three ways: ontologically, as the *limen* between presence and absence, or the dimension of deferred presence that encompasses language, technics, media, but also, as Chapter 6 discusses, can extend to all ontologies; epistemologically, as the *limen* between knowing and not-knowing, or the impossibility of fully knowing something, or of knowing something at all levels; and phenomenologically, as the traces we encounter in the world of this ontological and epistemological liminality. The dissertation's structure reflects this framing. Chapter 4 offers a reflection on AI and the epistemologically liminal, Chapter 5 continues into the ontologically liminal, Chapter 6 develops the tools I use to frame and query phenomena through the lens of spectrality, and Chapter 7 operationalises them.

The dissertation's prologue (Section 1.1) exemplifies the points discussed above, with the Voder anecdote foregrounding an interpretation of spectrality as discourse. Crucially, one in which the epistemologically and the ontologically liminal coexist and overlap. Indeed, as one listens to the audio excerpt, the informational bottleneck of the recording frees the characters from their bodily constraints, diffusing their presence through space and time. In this liminal space, the boundaries between human and non-human bodies blur (e.g. Harper receding as the Voder emerges and vice versa, or the Voder's uncertain and ever-changing ontological status), like ghostly actors on an immaterial stage. Presence becomes (or reveals itself as always already) performatively constructed.

Having established grounds firm enough to sustain the discussion, in the following sections I examine different contributions to this theme that I deem of particular relevance to the dissertation: first, in the humanities (Section 2.5.1); second, in music (Sections 2.5.2 and 2.5.3); and third, in AI artistic and critical research (Section 2.5.4).

2.5.1 *Hauntology*

The ideal point of departure for this overview is Derrida's *Specters of Marx* (1994a), as this monograph is considered the main catalyst for the shift in academic orientation wherein spectrality began to arise as a new and lively area of investigation (Blanco and Peeren 2013, 3).

The notion of *hauntology* that Derrida develops in *Specters of Marx* stems from his critique of Western thought's tendency to focus on presence and stability of meaning, or the

metaphysics of presence. More specifically, hauntology represents a consequential (and historicising) move from the thinker's treatment of language as inherently deferring, with any particular term acquiring its meaning, rather than from a positive quality, from its difference to other terms. Similarly, Derrida views all ontologies as intrinsically haunted by the negatives they exclude, and acquiring their substance by the effect of a performativity that conjures the *thing* (in the sense of the ghost) into presence "only in a movement of exorcism" (Derrida 1994a, 202).

Hauntology represents a crucial expansion of Derrida's deconstruction into the temporal and mnemonic domain, one in which presence and absence, difference and deferral, exist in a movement of relation to what is no longer or not yet (Hägglund 2008, 1). This expansion inevitably prompts a reflection on the spectralising processes introduced by technics and media, which according to Derrida require a thinking that supersedes ontological certainty (Derrida 1994a, 63).

This move, crucial for the dissertation's focus on AI, finds continuity in Bernard Stiegler's reflections around deferred presence in *Technics and Time*. Here, the technical-mnemonic apparatus of photography reveals "the *différance* at work in all presence" (Stiegler 2009, 20), an *ortho-thesis* wherein the perceptual evidence of what has been inevitably defers to what is not.¹⁵ Stiegler's example of Anita Ekberg watching herself in Fellini's *Intervista*, and remembering her younger self in Fellini's *La Dolce Vita* thirty years later, epitomises this, with Ekberg "return[ing] to herself as her own phantom" (Stiegler 2009, 22), her presence circularly deferred by the same technical apparatus that captures it.

The publication of *Specters of Marx* prompted in the subsequent years an exponential diffusion of the discourse on spectrality across disciplines, from media studies to sociology, from archaeology to ethnography.¹⁶ Sconce, for instance, investigates haunting in modern media, showing how contemporary notions of spectrality are ingrained in media technologies, and how the nature of this ingrainment shifted in the twentieth century towards textuality (Sconce 2005). Smith and Gordon instead approach haunting from a sociological and political perspective, arguing for a realism that includes spectres and ghosts, understood as estranging circumstances, gaps, and fractures in our understanding of the social world. Here, haunting

¹⁵ Stiegler writes: "a straightforward position, the upright honesty of the past, of the recording of what happens, of what happened, a right, true memory; and yet this exactitude unavoidably gives rise to an imprecision. This memory, which is only true as a mirror - is reflected in a symmetry by which it is reversed: in photographic specularity, the "right" is always already the "left," the true always already the false" (Stiegler 2009, 20).

¹⁶ For a thorough overview of spectrality's many areas of application, see Blanco and Peeren (2013).

reveals the “always incomplete forms of containment and repression ceaselessly directed toward us” (Smith and Gordon 1999, xvi). How these estranging circumstances manifest affectively and may thus be apprehended and analysed, is a crucial and open question, and one that Chapter 6 aims to thoroughly address.

Recent scholarship is also increasingly exploring the posthumanist dimension of spectrality, a dimension that the framework I offer in Chapter 6 adopts through Fisher’s own (critical) posthumanist stance. Sarah Rich applies hauntology theory to the study of shipwrecks, developing a speculative hauntographic practice (see Section 2.5.5 for a genealogy of the term *hauntography*) of “imagin[ing] the uncanny spatial and temporal ambiguities and tensions of a liminal object that is both present and absent” (Rich 2021, 14). This haunting from artefacts that have surpassed their original functions and are now entering into relation with the non-human world inevitably destabilises the dichotomies of human and non-human, life and death, presence and absence, nature and nurture.

A similarly posthumanist perspective is offered by Justin Armstrong, who reflects on the possibility of an ethnography of absence, where traces emerge through *ghost texts*, invisible dialogues between people and their material surroundings, resonances of “subjective and non-human presences and culture” (Armstrong 2010). For the purposes of this dissertation, it is relevant to note how these theories offer a particularly fertile ground for thinking of the human category as an unstable, even a haunted, construct. Chapters 6 and 7 will indeed include preliminary reflections on what kind of posthumanist world is the one inhabited by spectres.

2.5.2 *Spectrality in Music*

The theories discussed above encompass philosophy, media studies, cultural studies, ethnography. Yet, artistic practices and music in particular offer a particularly fertile ground to speculate around the epistemologically and ontologically liminal. Indeed, it may be argued that a certain spectrality accompanies the very emergence of electroacoustic music, since the technologies of reproduction, processing and transmission it makes use of grant an independent existence to sound, as if separating a shadow from the body that makes it.

In the fields of soundscape composition and ecoacoustics, R. Murray Schafer introduces the notion of *schizophonia* to describe how sounds are given an amplified and independent existence through their severing from their natural origins (Schafer 1969, 46). According to Schafer and the members of the Canadian World Soundscape Project (WSP), in

the new soundscapes brought about by modern technics, sounds are increasingly experienced in isolation from their sources (e.g., the “hole in the head” that produces a singing voice versus the song heard on the radio, as Schafer puts it (Schafer 1969, 44)) and detached from the bodies that make them; and, as modern life becomes ventriloquised, even from the gestures that shape them.

Schizophonia is endowed with spectral features by Schafer himself, who conjugates the notion through the legend of Dracula. As the story goes, the count returned to Romania from Paris, mourning the death of a famous opera singer with whom he had fallen in love. He brought back with him a white statue of the singer and her recordings, which he began to play at night in the solitude of his mansion through a phonograph placed close to the singer’s statue. In those evenings, the villagers working outside the count’s estate were terrified to hear the crackling sounds of the singer’s voice through the record player, and in glimpsing through the curtains the motionless shadow of a woman (Schafer 1969, 44).

In the context of the Parisian Groupe de Recherches Musicales (GRM) founded by Pierre Schaeffer, the separation of sound and source instead informs the emergence of an aesthetics of objectified sounds. As in the case of schizophonia, the term *acousmatic*, that famously identifies the practices of this group, originates from a separation between source and sound: the legend holds that Pythagoras’ initiated disciples, or *akousmatikoi*, would listen to the master speak from behind a veil to focus on *logos* and avoid the distractions of vision. Similarly, in acousmatic composition sounds are treated as phenomenologically independent acoustic formations to engage with through a practice of reduced listening, in which the source is deliberately removed from perception (Schaeffer 1998).

While theoretically clear-cut, in the practices of electroacoustic musicians both the critique of schizophrenic sounds and the acousmatic objectification of sound seem to represent two receding horizons (albeit different ones, the former rooted in ecological ethics and the latter in musical aesthetics). The evocative potential of sound, its tendency to stick to people, places, practices, memories, bodies and gestures, at all times pushes to expand the acousmatic work beyond the “here and now” of reduced listening. Similarly, the very practices from the WSP group inform a practical engagement with split sounds that seems at odds with their monolithic critique of reproduction technologies in use.

In a sense, however spectrally depicted, the schizophrenic split and the experience of reduced listening still pertain to the binaries of a metaphysics of presence, of which absence is nothing but the flip side. As Jonathan Sterne aptly puts it, there is an assumption that “at some time prior to the invention of sound-reproduction technologies, the body was whole,

undamaged, and phenomenologically coherent” (Sterne 2003, 21). Sound-reproduction technologies feature an inevitable commingling of presence and absence (and thus, to conjure Derrida, do not belong to ontology (Derrida 1994a, 63)), and even the idea that these reproduce an original on the other end is a modern assumption that responds to precise commercial and cultural reasons (Sterne 2003, 25).

An alternative view on schizophonia, one that does not exhibit Schafer’s “monolithic anxiety about the jeopardy to primal originality,” is offered by Steven Feld (Feld 1996), who from an ethnographic perspective re-contextualises schizophonia around practices located in the situations, flows, phases, and circulation patterns of sound recordings characterising how recordings move in and out of short-term and long-term commodity states. In this light, the schizophonically split sound should be reimagined as an inherently paradoxical phenomenon, in which the aura of authenticity can partially survive successive processes of re-presentation. This kind of spectrality inevitably brings the effects of mediation technologies to the fore, whose workings the artists discussed in the following section exploit creatively.

2.5.3 *Sonic Hauntology*

The spectral dynamics of reproduction and mediation technologies manifest in the musical counterculture as well, for instance, in dub’s *remixology* practice. Emerging from the Jamaican *sound system* tradition of the 1970s and later influencing UK rave culture, dub transforms existing recordings through extensive processes of reverberation and delay into ethereal soundscapes, where the original melodies and voices become traces of their former selves, barely audible behind the magnified drum and bass. This tradition, in itself quintessentially spectral in how it deliberately transforms and decontextualises the originals into something other and yet eerily familiar, is foundational to the practices of the inhomogeneous group of artists that Fisher describes as *sonic hauntologists* (Fisher 2013b). This phrase identifies a constellation of producers working in the United Kingdom at the end of the twentieth century, whose spectral aesthetics establish a resistance-through-culture against the neoliberal foreclosure of imaginable futures.

According to Fisher, artists such as Tricky, Goldie, the Caretaker, and William Basinski overlap temporally and semantically distant sonic signifiers to re-imagine the futures lost against neoliberalism’s naturalising claims, challenging the idea of historical time coming to an end as famously proclaimed by the elites of the Western intelligentsia (Fukuyama 1989).

Chapter 6 draws extensively on Fisher's sonic hauntology and other related works of his, which are placed in a dialogue with Derrida's *Specters of Marx* in order to develop the hauntographic framework.

2.5.4 Spectrality and AI

The exploration of spectrality as applied to AI that this dissertation offers is grounded on a small corpus of works describing generative technologies as spectral or even hauntological. Among these scholarly contributions, three are of particular relevance to the framework introduced in Chapter 6.

The first is Calvino's prescient and thus haunting reflection in *Cybernetics and Ghosts* (Calvino 1987). In this work, Calvino imagines cybernetic literature machines capable of generating texts through pure combinatorial processes. Drawing on oral storytelling, structural linguistics and information theory, Calvino views literature as a combinatorial game operating on a finite set of elements, whose permutations can bring about new narratives and ways of knowing. While literature machines are potentially capable of performing any possible permutation of a given material, the relevance of the permutation depends on its friction against human culture. In other words, the machine's ability to conjure ghosts is grounded (i) on its ability to unwittingly step out of cultural norms, and (ii) on the inevitable encounters with "the hidden ghosts of the individual and of his society" that this crossing entails. Little did Calvino know, though, that the cybernetic systems he was imagining would converge toward the statistical average of the knowledge whose boundaries they ought to expand, ultimately reinforcing the very biases and cultural norms that constitute these boundaries.

If not by design, the way to the truth of a system (whether a machine, the psyche or the entirety of human society) thus leads through what appears as a "pathological, marginal and accidental distortion of this system, [such as] slips of the tongue, dreams, symptoms, economic crises" (Žižek 2009, 101). Yair Rubinstein's framing of AI as an inherently hauntological technology follows this line of inquiry. The scholar postulates that similarly to how sonic hauntology's eeriness stems from how it uses new technology to "remediate older sonic artefacts in ways that deliberately upset how the past is usually represented," generative AI reassembles the timelines of the source material in novel, eerie ways, and, through temporal and causal disjunctures, re-enacts the contradictions hiding underneath capitalism's atemporalities (Rubinstein 2020). While for Calvino the "stepping outside" is a function of the

recombinatory potential of generative technologies, here it represents the result of their limitations in interpreting data in human terms. The modes of hauntographic analysis I describe in Chapter 6 draw on this perspective, while the insights from Chapter 7 attempt to further investigate the nature of these contradictions, for instance, through the notion of *spectrophonia* (see Section 7.4).

A third relevant work that this dissertation draws upon, and which I return to in multiple instances, is Jennifer Walshe’s essay *Ghosts of the Hidden Layer* (Walshe 2023), in which the composer explores the haunted nature of NAS systems in relation to their architectural opacity and to the evocative power of the human voice. Walshe discusses the training of AI models on recordings of her voice, and utters the question introduced in Section 1.2: “who is listening, and how?” This question represents one of the Ariadne’s threads in this dissertation. Walshe draws on Chion’s notion of *acousmêtre* (the omnipotent, disembodied presence that exists as long as it escapes the body producing it), wondering what the body behind an AI-generated voice would look like. The problem is that no-body is there, nor ever was, anchoring that voice to presence. *Mouja+*, the piece I discuss in Chapter 5, deals precisely with this impossible referentiality, but extends it beyond voice and to all AI-generated sonic signifiers.

2.5.5 A Genealogy of Hauntography

As a conclusion to this chapter, I would like to offer a brief overview of the different appearances of the term *hauntography*, which I chose to describe the framework offered herein.

The earliest appearance of the term *hauntography* I could trace dates back to February 2009, when a blog focusing on pop culture and music used it as a portmanteau for a series of posts and collaborative readings of M.R. James’ ghost stories.¹⁷

Sarah Rich (see Section 2.5.1) introduced the term in scholarly discourse as early as 2016, applying it specifically to shipwreck archaeology. Rich parallels hauntography with *ontography*, a methodology of Object Oriented Ontology (OOO) that Bogost frames as the “revelation of object relationships without necessarily offering clarification or description of

¹⁷ <https://freakytrigger.co.uk/ft/2009/02/hauntography-the-ghost-stories-of-m-r-james>. The term reappears years later in another UK blog focusing on horror narrative (hypnogoria.blogspot.com).

any kind,”¹⁸ (Bogost 2012, 38) describing a *hauntograph* as the “ontograph for the revenant” (Rich 2021, 14).¹⁹

Hauntography is also discussed in 2020 by Neil Gorman to describe a hauntology-informed approach to ethnography, where a researcher spends time among a group to investigate the spectres haunting it.²⁰ Finally, the term is applied to fine-art photography by Felicity Tsering Chödrön Hamer (2023), who frames it as the mapping of personally significant, absent photographic objects and the bereavement that arises from this, and by Mark Robinson, who describes it as an approach to producing work that is manifestly hauntological in its effect.²¹

As the above overview shows, there exist as many hauntographies as there are spectres, spanning narrative, media theory, ethnography, photography, and archaeology. This heterogeneity is at least in part attributable to the epistemological flexibility that the term affords, anchored as it is in the hauntological tradition, yet allowing the researcher to depart from it in substantial ways.

This is also where the appeal of the term *hauntography* lies for the purpose of this dissertation. As will become apparent in subsequent chapters, hauntography substantially draws on and at the same time departs in meaningful ways from the two main theorists of hauntology, namely Jacques Derrida and Mark Fisher, with the aim of operationalising their thinking into a framework that remains flexible and yet precise enough to suit different methodological applications, including my own practice with AI.

18 According to Bogost, examples of ontographs are verbal and visual lists, exploded views and ontographic machines such as video games mapping abstract gestures to encyclopedic accounts of things.

19 A similar parallel between hauntography and ontography is also traced by the blog hauntography.net.

20 <https://gorman.blot.im/hauntography>

21 <https://www.markrobinsonphotography.com/blog/dissertation>

3 Frameworks and Methods

3.1 Introduction

This chapter outlines the methodological approach that guides the research. I position the methodology within three macro-areas: practice research, which grounds the investigation by generating the instruments and practices that enable it; material semiotics, which enables the analysis of the discourse emerging from the instruments in use; and close reading, used to translate that discourse into the theoretical framework developed in Chapter 6 and to apply it to artist interviews. The chapter organises the discussion into three sections, followed by an ethics section.

The first section (3.3) situates the dissertation's methodological backbone within practice research. Definitions of practice-based and practice-led research vary considerably and sometimes conflict. Rather than fully adhering to one particular strand, I clarify how I understand practice research in this context, discussing how the dissertation's methodology both adheres to and diverges from common approaches.

The second section (3.4) describes the dissertation's primary mode of engagement with AI-based musical interfaces and the bodies that perform them as hermeneutic. From this stance, technologies and bodies are interpreted in their meaning-making capacity, with material semiotics providing the analytical framework to investigate their shifting relations.

The third section (3.5) focuses primarily on the methods applied in Chapters 6 and 7. Responding to findings from practice, these chapters engage with core literature on spectrality, and in particular with Jacques Derrida's and Mark Fisher's work, through comparative close reading methods. Having developed hauntography, the framework is then used to interrogate practice. Here, hauntography is applied primarily to the third-person accounts from practitioners I collaborated with, and secondarily to my artistic and design practice.

Before introducing the three methodological clusters and their specific methods, in the next section I discuss the dissertation's epistemological architecture and describe hauntography as a theoretical framework informed by and in turn informing the research.

3.2 Methods, Frameworks, Metaphors

Across the humanities and social sciences, frameworks typically function as conceptual structures, lenses, and anchors that ground and guide research. Methods, conversely, describe the specific procedures and techniques for gathering and analysing data. Methodologies, as the word implies, organise these methods according to an overarching rationale, or *logos*.

In this dissertation, frameworks and methods function in an iterative loop in which both change and adapt as the process unfolds rather than being monolithically established from the start. For instance, the *open work* (see Section 2.4) describes my initial epistemological position relative to instrument design and composition, but as Chapters 4 and 5 show, the insights from practice called for an extension of this framework. This expansion urged moving from artistic practice to theory (Chapter 6), and from there, again, validating the theory by observing how it allows one to interpret the practice (Chapter 7). This dynamic of interaction and iteration, of reciprocal destabilisation, at times even of messiness, is typical of artistic practice (see Section 3.3.2), and was crucial for the emergence of the insights discussed in this dissertation.

A relevant distinction in the literature separates conceptual frameworks from theoretical frameworks. For Kivunja (2018), theoretical frameworks define knowledge structures that guide data analysis. On the other hand, conceptual frameworks relate to all the concepts and ideas that occupy one's mind in contemplating, planning, implementing, and concluding a research project. Importantly, theoretical frameworks do not necessarily predate the research, and can involve a theory that is developed in the course of the dissertation study (Grant and Osanloo 2014).

These distinctions are sufficient to situate hauntography within the dissertation's epistemological arc as a theoretical framework, which informs but remains distinct from the methods I apply. I describe hauntography as a framework because it aims to offer a toolbox of intertwined notions that is flexible enough to adapt to different methodologies (from artistic practice to thematic analysis, from critical theory to design research) and subjects of investigation (thus applicable within and beyond music AI), cohesive enough to inform the perspective of the researcher yet non-prescriptive in terms of procedures and techniques. Furthermore, following Kivunja, it is worth clarifying that hauntography qualifies as a theoretical rather than conceptual framework in that it provides a knowledge structure that methods directly rely on (see for instance the deductive thematic analysis in Chapter 7).

3.2.1 Epistemological Metaphors

Conceptual and theoretical frameworks are often articulated from metaphors that synthesize their underlying epistemologies. Metaphors are core mechanisms of our conceptual system (Lakoff and Johnson 1980), which by highlighting some elements of a phenomenon necessarily obscure others, thus defining their epistemic accessibility (Möck 2022). For instance, the foundationalist claim that knowledge requires a ground or a base, Leibniz's notion of knowledge gained through the light of nature, or Peirce's parallel of knowledge assimilation to the interwoven fibres of a cable as opposed to the sequential links of a chain, are all metaphors that shape how we interpret reality (Thagard and Beam 2004).

This is even more evident in the context of music technology. As Impett (2023) argues, our engagement with music is shaped by *intuitive technologies*: metaphorical constructs derived from our technological environment and mediated experience, which function as sense-making operators. Rather than reflecting expert knowledge, these look more like plastic caricatures, reflecting the technological milieu in which one operates. Their malleability is precisely what makes them sense-making tools and endows them with epistemic capacity.

The metaphors introduced in Sections 2.4 and 2.5 function as intuitive technologies in precisely this sense. These modern accounts of openness and spectrality are evidently informed by precise technological milieus, the former by information theory, and the latter by media and recording technologies, and their epistemological value resides in their capacity to organise provisional, parallel sense-making processes as artistic practice takes place and is observed. Drawing on Eco's own words, in this dissertation I describe these as epistemological metaphors offering "a sort of transcendental scheme that allows us to comprehend new aspects of the world" (Eco 1989, 90). When these metaphors are articulated conceptually through Eco's degrees of openness or Derrida's hauntology, they expand into the dissertation's conceptual framework, which distils through Chapters 4 and 6 into hauntography's theoretical framework, and in turn informs the methods applied in the latter part of this dissertation.

Having established the dissertation's epistemological arc and sustained it through scholarly literature, the following sections outline the three methodological clusters that informed hauntography's emergence and application, starting from practice research.

3.3 Practice Research

Bolt and Barrett describe research in which artistic practice functions as the primary mode of investigation and knowledge production with the phrase *practice as research* (Bolt and Barrett 2019). Within this broad field, some scholars distinguish two macro-areas of artistic investigation, practice-led and practice-based research, with definitions that vary considerably. For instance, Gray (1996) describes practice-led research as research which is initiated in practice and in which the research strategy is carried out through practice and applies methods that practitioners are familiar with. Candy (2006), on the other hand, defines practice-led research as that which leads primarily to new understandings about practice, distinguishing it from practice-based research, in which the artefact in itself represents the contribution to knowledge.

More broadly, for Smith and Dean (2009), practice-led research identifies (i) the work of art as a form of research and (ii) the ability of the work to generate research insights which might then be documented, theorised and generalised. Finally, Haseman views practice-led research as the foundation for a whole new research paradigm (performative research), entirely distinct from the qualitative methods of practice-based research in that contributions to knowledge are expressed in symbolic forms rather than text, and the research does not necessarily entail an initial problem formulation (Haseman 2006, 2, 5).

Among these definitions, this dissertation's methodology resonates most with Smith and Dean's formulation, in that it approaches artistic practice as research and theorises documented, generalisable insights that extend beyond practice itself. Nonetheless, I concur with Bulley and Şahin, who, drawing on Michael Biggs's work, argue that territorial debates between practice-led and practice-based research are ultimately unproductive (Bulley and Şahin 2021, 25). Rather than defending conceptual boundaries, what matters is articulating how specific methodological approaches apply to one's work. From this stance, the following subsections discuss how key aspects that scholars identify as characteristic of practice as research (subjectivity, emergence, circularity, and interdisciplinarity) are addressed in this research.²²

²² Here I use the phrase "practice as research" purposely, as a reference to Bolt and Barrett's influential work (2019), which under this phrase groups diverse contributions discussing practice-based, practice-led, performative research, and other sub-fields.

3.3.1 *Subjectivity*

Practice research acknowledges and values the researcher's embodied, first-person perspective as a legitimate epistemological position (Gray 1996, 13), and through such perspective foregrounds intuitive, embodied, tacit, imaginative and affective ways of knowing (Bulley and Şahin 2021, 38). As Polanyi puts it in describing tacit knowledge (and the first-person investigation on XAI in Section 4.5 exemplifies), “we can know more than we can tell” (1966, 4). By paying close attention to situated relations between individual subjectivities and objective phenomena we can thus begin to reveal and articulate this tacit knowledge (Bolt and Barrett 2019, 4).

Chapters 4 and 5 present first-person reports and reflections on my artistic engagement with AI, which I captured as they unfolded through constant journaling. This documentation proved essential in revealing discursive patterns and metaphors as they surfaced in collaborations and performances, and to reflect on the affective dimension that accompanied them. Throughout the dissertation, I report episodes where metaphors of haunting and spectrality emerged organically, and which I captured through journaling: for instance, my colleague Jack Armitage's spontaneous Ouija board comparison when encountering my instrument Stacco for the first time (see Section 4.7.1), or composer Robert Laidlow's idea of using the same instrument for a piece reflecting on one's communication with the ghosts of a past self (see Section 7.2.2).

3.3.2 *Emergence*

In discussing practice-led research, Gray shows how practitioners immerse themselves in their practice and through this engagement new problems and questions arise (Gray 1996, 3). Relatedly, Haseman discusses how researchers often do not begin with defined questions, but rather from an “enthusiasm of practice” from which research questions emerge (Haseman 2006, 3). This process is non-linear and “messy” (Denzin and Lincoln 2003, 7), and inevitably entails false starts and research pivots that practitioners make sense of only in retrospect.

This dissertation differs from the purely emergent stance described above, in that it began with a defined project and precise research questions around XAI. At the same time, I deeply enjoyed the freedom to experiment with materials and processes, to engage with the questions from an enthusiastic and sometimes even messy stance, and to rearrange my

questions as findings from practice opened up new research trajectories. I let this “messiness” intentionally show through in Chapter 4, where I present the reader with the entangled knot of theories, materials, studies, and reflections I immersed myself in during the first year of the project. From this knot, a spectral thread begins to unravel, conducting the reader, this time quite linearly, through the haunted work (Chapter 5), hauntology theory (Chapter 6), and hauntography in practice (Chapter 7).

3.3.3 *Circularity*

Practice research typically operates through iterative cycles in which practice generates theoretical insights that then return to inform and interrogate practice. Smith and Dean describe this dual iterative movement as *practice-led research* and *research-led practice*. These dynamics should not be viewed as separate, but rather as “interwoven in an iterative cyclic web” (Smith and Dean 2009, 2). Similarly, in Bolt and Barrett (2019, 152), Haseman reports on theatre director and PhD candidate David Fenton, whose research method was grounded in a cycle of workshops, reflections, and rehearsals which he describes as a double-loop of reflection-in-action and reflection-on-action.

This dissertation applies such circularity across multiple scales. At the micro-level, each design iteration involved cycles of making, performing, reflecting, and redesigning: I developed two versions of my interface Thales (see Section 4.3) and three iterations of Stacco (see Section 4.4), with each version responding to new insights and needs emerging from use. Similarly, I constructed my pieces and performances iteratively: *Mouja* (Section 5.2) took shape over dozens of performances distributed in three years, with *Mouja+* (Section 5.4) representing the most expanded iteration and incorporating new questions and methods emerging from earlier performances.

At the macro-level, the dissertation’s structure in itself enacts a certain circularity, with the practice documented in Chapters 4 and 5 provoking new reflections and questions, Chapter 6 translating these into hauntography’s theoretical framework, and Chapter 7 returning to observe practice through it. Crucially, this final phase expands this circle of creation and reflection beyond my subjective experience to include the perspectives of six artists I collaborated with, thus addressing the need for replicability and transferability of research insights in practice research evidenced by Bolt and Barrett (2019, 13).

Finally, this circularity of practice and theory extends into hauntography itself, whose vocabulary applies to creative approaches, with the categories of the spectral viewed as fictional modes, and to analytical ones, with the categories of the spectral viewed as affective modes through which one apprehends spectral phenomena in the world.

3.3.4 Interdisciplinarity

One aspect that scholars discussing practice research in its different formulations (e.g. practice research, practice-based research, performative research, etc.) tend to agree upon is the inherently interdisciplinary nature of the methods it draws upon, with practitioners drawing from multiple established fields (e.g. reflective practice, participant observation, performance ethnography, ethnodrama, biographical/autobiographical/narrative inquiry (Haseman 2006, 3)). Such an inclusive and diverse approach allows artists to validate the methods and practices they are most familiar with (Gray 1996, 3) and new ones that might be designed throughout the process (Haseman 2006, 3).

Such interdisciplinarity characterises this dissertation as well. From the outset, I developed my practice around design and artistic practice, gaining new insights through user studies (see Section 4.6), and evaluating the process through reflective journaling. As practice and reflection progressed, themes began to emerge and naturally necessitated theoretical reflection, which I approached through close reading of texts as detailed in Section 3.5. The following sections discuss these different methods, beginning with material semiotics.

3.4 Material Semiotics

NIME design necessarily involves the materialisation of worldviews, musical representations, and cultural values. As Giacomo Lepri argues, “while digital music communities are sometimes criticised for being preoccupied with technical factors, an equal challenge is not the lack of aesthetic reflection but rather a surfeit of latent cultural assumptions” (Lepri 2022, 20). This insight might be true for all (musical) instruments, but it is relevant here to consider Magnusson’s argument that while acoustic instruments are imbued with knowledge, their digital counterparts are imbued with an even higher level of epistemic dependencies

(Magnusson 2019, 57). DMIs may thus be viewed as epistemic tools, endowed with such a “high degree of symbolic pertinence that [they] become a system of knowledge and thinking in [their] own terms” (Magnusson 2009).

This increased epistemic density entails a certain responsibility for DMI designers, who, whether deliberately or not, inevitably end up inscribing certain aesthetics, assumptions, social and cultural values onto them. These, then, propagate in time and space, and resurface anew through situated negotiations with the instrument-in-use. Crucially, this is not merely a process of representation of old norms and values, but rather of remediation and renegotiation through the particular affordances and constraints of the interface,²³ which thus becomes what Scolari describes as a *sense production device* (Scolari 2009). To analyse how these values and assumptions are materially inscribed in DMIs, and how they produce meaning in their use, I turn to material semiotics.

Material semiotics describes an interdisciplinary and diverse constellation of approaches to analysing the social world, in which the material domain inherently carries, shapes, and is shaped by meaning. More precisely, Law defines it as a “set of tools and sensibilities for exploring how practices in the social world [...] form weaves that are simultaneously semiotic (because they are relational, and/or they carry meanings) and material (because they are about the physical stuff caught up and shaped in those relations)” (Law 2023). A question lingers in this definition regarding the ontological status of the materials caught up in this web: should these be viewed as mediators, as “weak actors,” or as actors on par with the humans they entangle with?

Within material semiotics, Actor-Network Theory (ANT) proposes what might arguably be described as a flat ontology of situated translations between human and non-human actants (Latour 2005). As Pickering argues in interpreting Callon and Latour (1992), semiotics is crucial to ANT in that it enables thinking symmetrically about human and non-human agents, since in texts agents, actors and actants “are continually coming into being, fading away, moving around, changing places with one another, and so on” (Pickering 1995, 12), something that the discursive coming and going of Helen Harper and the Voder in the dissertation’s prologue exemplifies.

While I maintain some reservations about ANT’s tendency to flatten power dynamics and accountabilities, I recognise that the non-human world contributes to meaning-making

23 Building on Gibson’s theory of affordances, Magnusson describes how affordances point to features that make things possible, whereas constraints define the limits of the possible (Magnusson 2010).

more often than anthropocentric positions tend to acknowledge. In this dissertation, I do not use semiotics to advance claims about the ontological status of interfaces or AI systems. Rather, quite hauntologically, I am interested in inhabiting their ontological ambiguities and contradictions as productive loci.

More than ANT, the above epistemological positioning echoes Donna Haraway's doctrine of embodied feminist objectivity and situated knowledges, which insists on material and semiotic linkages between meaning, technologies, and bodies, while maintaining a critical stance toward normative ways of knowing (Haraway 1988). Rather than flattening ontologies, this view allows conflicting perspectives to coexist and contribute to meaning-making. This space of contradiction, in particular as far as presence and absence (as being here-and-now, but also being *tout-court*) are concerned, is something that hauntography, and the hauntology from which it derives, attempts to preserve as an epistemological precondition.

3.4.1 *Material Semiotics in Practice*

For the analytical task of examining how DMIs inscribe and communicate meaning, Scolari provides operational categories that neither require nor foreclose ontological commitments. Drawing upon Eco's work on the relevance of the reader's role in interpreting texts, and viewing "the ambiguous game between signification and interpretation played by designers and users [...] as a cooperative one (designers ↔ interface ↔ user)," he distinguishes four research levels in the semiotic analysis of HCI interfaces at work (Scolari 2009):

Plastic Level, at which the researcher identifies basic elements such as lines, shapes, colours, textures, positions, and investigates how these contribute to sense production.

Figurative Level, relative to the representative scene, which is made up of objects, areas and mutations.

Communicative Level, relative to the research into the spatial location and narrative in relation to the scene.

Meta-communicative Level, relative to the analysis of the position of the receiver in relation to the entire communicative context in which the interface is deployed.

Drawing on these semiotic levels, Chapter 4 observes emergent meaning at the plastic and figurative levels as applied to the two novel NIMes I introduce. This is achieved in two ways: first, by reflecting on how their physical shapes informed the embodied interpretation of the AI model at work and of the organisation of the latent space (see Section 4.5); second, by discussing how the affordances of the interfaces informed the emergence of a particular cluster of metaphors (see Section 4.7). Chapter 5, instead, investigates the interfaces-in-use within different performative scenarios, observing how these inform the communicative level of the performance planning and organisation, and finally the meta-communicative level of the performance as interpreted by the viewer.

3.4.2 A Hermeneutic Approach to Embodiment

To expand the reflection into the communicative and meta-communicative levels means inserting the interfaces within the much broader social webs of artistic practice, and thus in relation to the people who perform them. Through practice, one incorporates instruments into the bodily schema, creating entangled body-instrument assemblages where boundaries inevitably blur (De Souza 2017). These dynamics of shifting relations between tools and human bodies are a relevant object of study within phenomenology.

One of the first accounts of embodiment relations may be found in Heidegger's notions of *Zuhandenheit* and *Vorhandenheit* (*ready-to-hand* and *present-at-hand*) (Brey 2000). One exists in a ready-to-hand relation with a technical object, say, a hammer, until it breaks, and this breaking compels a distancing and its examination as "other," at which point it becomes present-at-hand and distinct from the self. Merleau-Ponty extends Heidegger's insight by reframing the tool's transparency in ready-to-hand relation as bodily potentiality. Here, the body itself becomes a field of interaction with the environment, and its parts are understood, rather than as spatial relations, through their capacity to enter into one's actions and projects (Merleau-Ponty 1990).

While Heidegger and Merleau-Ponty focus on understanding perception and the prosthetic extension of the human body, Don Ihde's postphenomenological approach argues for a phenomenology of human-technology relations in which technologies actively mediate and transform our experience of the world, and are thus endowed with technological intentionality. Ihde distinguishes four types of human-technology relations: embodiment,

where the technology withdraws from perception while we experience the world through it, thus becoming a partially transparent extension of the body (for instance, a pair of eyeglasses); hermeneutic, where technology is interpreted to interpret the world (for instance, a thermometer); alterity, where technology becomes “other” from us; background, where relation happens through the technology’s embeddedness in a certain context (Ihde 1996).

Musical instruments can exemplify all of these modes. Indeed, when one selects a certain technology, one inevitably chooses a lens, a background theory, a specific worldview, and thus defines the interpretative scope of potential outcomes including the conceptual resolution and type of embodiment (Magnusson 2019, 51). Yet technologies do not simply implement pre-conceived intentions. Once in use, they resist, surprise, and reconfigure, through what Pickering terms the *mangle of practice*, in which human goals and material constraints co-constitute each other through negotiation (Pickering 1995). According to this view, researchers set certain goals and design experiments, but on the other hand materials resist, reconfigure unpredictably, and push back, thus forcing accommodation and adjustment. Pickering’s metaphor describes with a certain accuracy my understanding of learning, teaching, and playing an instrument, both at the macro and microscale: what an external viewer might describe as mastery feels from the inside as a constant act of negotiation, a “tightrope walking” between continuous states of imbalance that thanks to practice appear as a seamless flow.

Throughout the dissertation, the reader might discern many types of embodied relations and material resistance, such as embodied (e.g. Section 4.5), hermeneutic (e.g. Section 4.7) and alterity (e.g. Section 5.2) relations. While all these relations are evidenced and discussed in different ways, this dissertation adopts a hermeneutic approach to embodiment, in the sense that it observes and interprets the instruments in use to understand how shifting body-instrument relations translate discursively on and beyond the stage, and how these discursive practices contribute to making sense of intelligent technologies.

This does not imply a denial of the very existence of bodies as material entities, but rather, within what Violi terms a radical constructivist perspective, the recognition that bodies in-themselves are negotiated and co-constructed, that there exist as many types of bodies as there exist discourses around them, that some of these are present despite their material counterparts (as ghosts teach us, I would argue), and that “we can only reach these bodies through [...] semiosis” (Violi 2008).

3.5 Close Reading of Texts

The third methodological cluster of this dissertation may be described as close reading of texts. Close reading implies a careful and sustained investment in the interpretation of a text, with attention to discourse and formal structure, but also on individual words and syntactic choices. This is different, for instance, from distant reading or surface reading techniques, in which the researcher engages less deeply with a certain work, but expands the analysis into a much larger number of texts in order to identify emerging themes.

3.5.1 *Comparative Close Reading*

Chapter 6 establishes the hauntographic framework through comparative close reading of two core texts. These are Jacques Derrida's *Specters of Marx* (Derrida 1994a), which is seminal in establishing the notion of hauntology, and Mark Fisher's *The Weird and the Eerie* (Fisher 2016), which does not directly discuss spectrality or hauntology theory, but which I situate in continuity with it. The comparative reading of these texts intends to show how the latter work operationalises some crucial features of spectral phenomena that in Derrida remain undefined. To achieve this goal, I ground the analysis on three methodological axes that the following paragraphs describe and motivate through existing literature.

Matthew Arnold, who coined the term *comparative literature* in 1857, argued that “no single event, no single literature is adequately comprehended except in relation to other events, to other literature” (Brown 2013). Yet, even when we do not explicitly bounce a text against some other texts that we read, we inevitably compare a text with the memory or the expectation we have of it. The specificity of comparative methods, then, is that these make the comparison explicit and analyse the modes of said comparison. Brown (2013) suggests a distinction of particular relevance to elucidate the methodology applied in Chapter 6, arguing that comparison may be approached either symmetrically or asymmetrically. In the case of the latter approach, a different quantity or quality of attention is dedicated to each of the two *comparanda*, making asymmetric comparative reading similar to much criticism which is not usually considered comparative. Symmetrical comparative analysis, which is the strategy I apply in Chapter 6, tends instead to pay a similar quantity and quality of attention to a discrete (in this case two) number of comparanda, that are thus not viewed hierarchically.

Another aspect that the comparative analysis applied in Chapter 6 acknowledges and treats with attention is the historical and cultural situatedness of the works. As Hawas notes in reflecting on English literary criticism's tendency to downplay these aspects, context and positionality are crucial in comparative work, and often involve a certain authorial political engagement (Hawas 2025). Chapter 6 extensively discusses the two main texts in relation to the historical context they stem from and culture they engage with, also by involving adjacent contemporary authors (for instance, Fredric Jameson, Jean-François Lyotard, Francis Fukuyama, Giorgio Agamben) that are useful to further clarify Derrida's and Fisher's positioning.

Finally, a third methodological consideration that is consequential to the question of historical and cultural situatedness refers to the need to establish a solid and shared thematic ground between the texts under scrutiny, so that these can, so to say, speak the same language. This is a typical approach in comparative works, and a necessary one in order for the researcher to effectively synthesise their similarities and differences. The comparative reading I propose establishes this common ground on different levels. First, as argued before, at a historical one, situating Derrida and Fisher in temporal continuity with each other and in the immediate aftermath of the political turmoil between Western and Eastern blocs at the end of the twentieth century; second, at the discursive level, where the authors ground their respective hauntologies on a shared critique of the rhetoric of "the end of history," which Derrida engages with through a critique of messianisms and Fisher reorients from the perspective of cultural production.

3.5.2 Comparative Close Reading in Practice

The dissertation operationalises the three methodological axes described above (symmetrical reading, situatedness, common thematic ground) as follows. I began by reading the texts and re-reading complex passages (in particular in Derrida), taking notes, summarising them, and iteratively going back to reading specific sections, until I developed a reasonably good understanding of each text. Both authors engage thoroughly with their contemporaries, but also with previous literature, for instance, Marx in the case of Derrida, and Jameson in the case of Fisher. I thus engaged with this secondary literature where necessary to the comprehension of the texts, and through this developed a more situated understanding of the primary sources. I then went back to the notes I took on the text, and read the two texts in dialogue with each other, looking for common thematic and semantic patterns.

For instance, throughout the analysis it became apparent how Derrida and Fisher share a certain interest in the fracturing and folding of historical time in relation to ideology, which they nevertheless articulate differently. This common interest determines their clustering around reflections on how the past, but also the future, haunt the stability of what we take as a given and consider almost as “natural.” This leads to the authors’ adoption of contiguous or shared terms and expressions, of which the phrase *out of joint* is perhaps the most evident one (see Section 6.5.1).

As the research progressed and a shared vocabulary and themes emerged, to maintain a certain critical distance from both authors I simultaneously engaged with literature that responds critically to Derrida’s and Fisher’s use of Marx.²⁴ This also allowed me to develop a critical stance toward their work, and from here, around the discourse on spectrality in music AI and beyond. This was crucial in developing hauntography, which inevitably reflects, along with the positionalities of the authors I analyse, my own positionality shaped by the commitment to practice and literature.

Having established and mapped the common vocabulary and themes, and having situated the authors along with their texts and positioned myself in relation to them, I began looking at where and how they diverge from each other. For instance, Fisher is highly critical of Derrida’s tendency to recede into epistemological indeterminacy, and frames the *weird* and the *eerie* as categories of the *strange* rather than of the *spectral*. These gaps were highly productive, as hauntography began to emerge through initial attempts at filling these interstices.

Here is precisely where the experiences and knowledge I gathered through practice, my personal understanding informed by the research and, inevitably, explicit and implicit assumptions, enter the picture. When faced with irreconcilable positions, I had to take sides, or discard both based on a certain rationale. I specify and motivate these rationales in Chapter 6, but broadly speaking, these are informed by practice (i.e. how well does this notion represent the phenomenon I experienced in my practice?), by my critical positioning in relation to the texts, or by the broader gap hauntography intends to address (i.e. does this notion inform and facilitate an approach to phenomena as spectral? Does it contribute to resolving the epistemological ambiguity that is limiting my application of spectral theories?). This led me to adopt some but not all of Derrida’s stances (for instance, I adopt the idea of the ghost as unknowable potential predating any differentiation, but discard his politically ambiguous use

²⁴ Of particular use was Domenico Losurdo’s *Western Marxism*, which traces and criticises the material evolution of the Western academic strand of Marxism (Losurdo 2017).

of the term *spirit*), some but not all of Fisher's positions (for instance, I adopt Fisher's critique of anthropomorphisms, but do not view the *weird* and the *eerie* as categories of the *strange*), and some entirely original claims justified by practice and engagement with the sources (for instance, the argument that the *weird* and the *eerie* are co-constitutive).

Having established a vocabulary that in Chapter 6 I cluster around the hauntographic theoretical framework, and traced the connections between the different terms within this vocabulary, in Chapter 7 I operationalise the framework through a thematic analysis of interviews with artists I collaborated with in the past three years, also informed by my practice and by discussions with lab members and visitors.

3.5.3 Hauntographic Analysis of Collaborative Practice

Chapter 7 applies deductive thematic analysis to interviews with six artists I collaborated with during my research at the Intelligent Instruments Lab. This method is also described as *theoretical thematic analysis*, in that a pre-existing theory or framework (in this case, hauntography) guides the coding and theme development. This is diametrically different from inductive thematic analysis, which starts with the data and from there derives the theory.

The six interviews involve artists whom my colleagues at the Intelligent Instruments Lab and I collaborated extensively with. These collaborations span the four years of the Intelligent Instruments Lab's activity, and encompass a group of international and local artists interested in incorporating AI technologies in their practices. Of the six interviews, five were carried out online, and one in person. All the interviews were recorded and transcribed via a local audio transcription tool.²⁵ Among all the collaborations, I selected this group of artists based on the diversity of their projects (both in terms of tools and processes), on their unawareness of my theoretical work, and on a certain tendency of their works to exhibit hauntographic features.

The interviews were semi-structured. Each interview was divided into two consecutive phases, a first one consisting of closed questions and a second, open-ended one. The closed phase consisted of six questions designed to frame the ongoing collaboration with the lab from the perspective of the participant, and to establish a biographical and artistic ground for the open-ended discussion. Questions were: (i) where are you from, and what is your musical

25 The Notes app on my local Apple machine.

background? (ii) could you describe your collaboration with the lab? (iii) what were (or will be) the outcomes from this collaboration? (iv) could you describe the role of AI in these works? (v) how does it feel to work with AI? (vi) how does working with AI compare with other methods you have experimented with? These questions were deliberately broad, with the hauntographic analysis taking place in the following open-ended phase, in which themes raised by the participants were investigated in more depth.

The thematic analysis of the transcriptions draws on a distinction developed in Chapter 6 between two categories, the *weird* and the *eerie*, that Section 6.5 frames as both modes of fiction and affective categories of the encounter with spectral phenomena. In applying hauntography analytically, the affective mode is of relevance. The thematic analysis proceeded through successive close readings of the transcribed text, followed by the coding of themes and keywords that refer to the *weird* and the *eerie* as described in Chapter 6, with the iterations lasting until no more hauntographic themes emerged. For instance, themes referring to changes in how the flow of time is perceived as reversed, fractured, non-existent, would be easily traced to the *weird*, whereas themes around disembodied agency, projection of presence, failure of absence, would refer to the *eerie*.

Since hauntography frames the *weird* and the *eerie* as co-constitutive, when a theme pertinent to one category was identified, it was coded and further investigated in the text from the perspective of the other category. Once the themes were identified, they were further analysed in the context of the single interview by using the other theoretical tools offered by hauntography (see Section 6.2) along with the sub-categories of the *weird* and the *eerie* identified by Fisher and summarised in Sections 6.5.1 and 6.5.2. These are then discussed in relation to the insights from practice offered in Chapters 4 and 5, and finally grouped together into the three thematic areas of the concluding discussion, namely AI NIMEs, AI and sound mediation, and AI in commodity form.

3.6 Ethics

This section covers the ethics research protocols. Ethical challenges arising from the use of AI more broadly and as part of this research are instead discussed throughout the dissertation: Sections 5.3 and 5.5.3 investigate the possibility of using dataset collection as a collaborative process to foster community kinship. Section 5.4.6 builds on the insights from performance

practice with a text-to-audio model to argue that open access to training data is essential to address issues regarding bias, data ownership and accountability, but also to facilitate informed interaction with generative systems. Section 7.4.2 offers a preliminary reflection on the application of the principles of soundscape composition to music AI. Finally, Section 7.5 expands from music AI to AI more broadly, showing how current narratives underpin its deployment as an extractive instrument.

This dissertation presents two studies involving artists from the Intelligent Instruments Lab's extended community. The first is discussed in Section 4.6, and the other in Section 7.2. In both cases, participants gave written informed consent to the interview, to their data being analysed for research purposes, and to their identities being revealed as in Section 7.2 or anonymised as in Section 4.6. The consent forms were defined as per the Intelligent Instruments Lab's commitments to ethics, which were approved by the Icelandic Ethics Research Committee at the outset of the research.

The same procedure applied in the training of models on data from third parties, such as in the case of the *End of Empire* installation in Section 5.3, in which written consent to the use of recordings was given. The installation was sponsored by a local institution of which the participants were members, focusing on social integration through art, and represents an entity that the community trusts and relies upon.²⁶ After the installation, the datasets along with the NAS models I trained were erased as previously agreed with the contributors.

Because the instruments I designed for my artistic research or in collaboration with artists typically feature NAS models, I had to curate or assist in curating many datasets of sounds. The datasets used for NAS in this dissertation are either open-source, or have been directly collected by the artists who used the models from their own recordings and with the consent of all the parties involved. Regarding personal datasets, models were deposited in a dedicated repository solely upon receipt of the data owner's written permission.²⁷

Finally, the NIMes I developed are all open-source, and their schematics, 3D models, printing instructions, components, software, and NAS models are all available on the Intelligent Instruments Lab's GitHub page.²⁸ In some instances, I developed tools based on MaxMSP, which is proprietary software. While MaxMSP patches can be freely downloaded and performed without a licence, to change them it is nevertheless necessary to purchase one. It is

26 <https://www.in-situ.org.uk/>

27 <https://huggingface.co/Intelligent-Instruments-Lab/rave-models>

28 <https://github.com/intelligent-instruments-lab>

my intention, in the coming months, to develop fully open-source, Pure Data-based versions of the same patches.

4 Inherent Scores

*One has to know. One has to know it.
One has to have knowledge.
Jacques Derrida (1994).*

4.1 Introduction

The aim of this chapter is twofold: first, to show how, due to initial findings emerging from practice, the research focus shifted from XAI to Un-XAI (see Section 2.3.5); second, to begin addressing the first research question through an exploration of the meaning-making capacity of what resists explanation. This latter process unfolds throughout Chapters 4 and 5. It progresses from the analysis of the plastic and figurative semiotic levels of the interfaces (these are discussed respectively in Sections 4.5 and 4.7), and in Chapter 5 expands into the communicative and meta-communicative levels (see Section 3.4.1).

Chapter 4 introduces a heterogeneous mix of instrument designs, user studies and artistic practice, which I divide into three sections: first, one introducing the new interfaces for musical expression (NIMEs) I developed at the outset of the research, and the rationale behind their design; second, a section in which I use the interfaces as part of my initial aim of investigating XAI through a pragmatic and embodied perspective (see Sections 2.3.3 and 2.3.4); third, a section in which, based on the insights from practice and user studies, the research project is reframed around the sense-making potential of what resists explanation, as per the first research question detailed in Section 1.3.

The first part (Sections 4.2 to 4.4) discusses inherent scores, open notation systems and musical interfaces wherein inscription and interface merge fluidly, and relates them to the affordances of latent navigation in neural audio synthesis (NAS) through their common exploratory character. Building on this, the chapter introduces two novel NIMEs, Thales (Figure 4.1) and Stacco (Figure 4.5). I describe the interfaces in use as *magnetic scores*, a type of inherent score where the musical inscription emerges relationally through magnetic interactions (Privato et al. 2023a). I designed Thales and Stacco to investigate and possibly accommodate the compositional and performative challenges posed by NAS, particularly the

apparently *arbitrary encoding*, *feature entanglement*, and *dimensional entanglement* that characterise latent space navigation (see Section 4.2.2).

The second part of the chapter (Sections 4.5 and 4.6) documents two studies in which Thales and Stacco are used to investigate XAI questions with real-time NAS models. The first consists of an investigation through practice on whether and how the multichannel spatial projection of sound might facilitate the learning of latent multidimensional topologies; the second uses Stacco to investigate whether and how a regularised latent space is preferred by artists over a non-regularised one when transferred onto another model. These investigations show that the idiosyncrasies of NAS models that XAI approaches seek to explain (a term that in Section 4.7 I literalise as the “making plain” of high-dimensional, entangled topologies) are often viewed by artists as creative constraints.

In the third part (Sections 4.7 and 4.8) I discuss how findings from the investigations described above prompted a reorientation of the research. Rather than looking at AI’s opacity as an inherent limitation to solve,²⁹ I began to consider it as a creative resource to be inhabited and explored. Hauntologically speaking, this shift represents a transition from attempting to *exorc-analyse* the ghosts-in-the-machine towards giving speech to them, as Derrida invites the scholar of the future to do (Derrida 1994a).³⁰ Starting from this theoretical perspective, I thus begin to articulate how the unfathomed in AI manifests aesthetically and perceptually as a question of presence and absence, setting the stage for the artistic investigation offered in Chapter 5 and for the hauntographic framework that will emerge in subsequent chapters. Employing the definition of spectrality as ontological and epistemological liminality presented in Section 2.5, this chapter primarily addresses the epistemological side, while Chapter 5 focuses on ontological liminality.

Before proceeding, I would like to emphasise how the content of this chapter, more than any of the successive ones, emerged from an enthusiasm rooted in practice guiding a reflexive process. As discussed in Section 3.3.2, research methodologies and questions in artistic research are often emergent and subject to repeated adjustment, and tend to shift throughout one’s enquiry. They thus involve a certain amount of creative messiness (Denzin and Lincoln 2003, 7). Through a mix of design and artistic practice, user studies and theoretical

29 Which certainly it can be, but not in all cases and especially not always in artistic practice, as shown in Section 4.6.

30 Derrida uses this term in *Specters of Marx* (*exorcizanalysis* in 1994b, p. 37; *exorc-analyze* in 1994a, p. 58). This notion may be understood in light of his critique of the metaphysics of presence, showing how the analysis of a phenomenon often hides the need to exorcise away the uncertainty surrounding it, to localise it once and for all.

reflection, the following sections introduce the messy tangle of forces (the ones, theoretical as well as material, that I could discern), from which the thread I follow in the subsequent chapters unfolds. These forces are not always crucial to make my points. Nonetheless, they are crucial to show where my points come from.³¹ As will become apparent through the reading of this dissertation, to exclude these forces in favour of a retrospectively constructed linear narrative would perform the same exorcisms that I seek to critique.

4.2 Background

4.2.1 *Inherent Scores*

Musical instruments function as tools for sound production but in that functioning also serve as mnemonic supports, materially inscribing the practices and theoretical systems that shape them as concretised forms of music theory (Magnusson 2019). This implies a certain overlay between the instrument and the notation, which in Digital Musical Instruments (DMIs) becomes even more evident and pronounced (Magnusson 2009). The notion of *composed instruments* epitomises this: it postulates a decoupling between sound and gesture, and the interweaving of the instrument with the musical work (Schnell and Battier 2002), thus extending the notational space beyond the physical constraints of the score, and incorporating all the structural, tangible and virtual components that support the inscription.

As a probe to investigate the merging between the instrument and the score, Tomás and Kaltenbrunner propose the notion of *inherent scores* (2014), tracing the origin of this concept back to Alvin Lucier, who, in describing the practices of the Sonic Art Union, stated that the scores were “inherent in the circuitry” developed by the members of the collective (Lucier 1998). This merging can be traced even further back in time, at least to 1967, when Gordon Mumma discussed how “designing and building of circuits is really composing” (Nyman 1974, 91). But even though Mumma and Lucier expressed the idea in plain language, a certain overlap between the score and the instrument has been clearly central to the practices of many other

³¹ As an example, in Section 4.3.1 I discuss how Thales was developed by literalising Barad’s principle of *intra-action* (Barad 2014). While in a few instances I do refer to Barad’s agential realism, this theoretical framework is not crucial to the discussion, but rather serves the purposes of documenting the unfolding of practice.

contemporary artists (David Tudor, Morton Subotnick, Suzanne Ciani, Éliane Radigue, Erkki Kurenniemi, to name a few) interested in rethinking the loci of production and inscription of the musical experience.

These types of notations introduce openness through non-linear navigation and discovery, endowing the score with an intrinsically exploratory character, and in this sense operationalise Eco’s notion of openness of the second degree (see Section 2.4). The following sections examine how this exploratory dimension may be leveraged to facilitate notation and performance with NAS systems. To begin this investigation, it is first necessary to establish what new challenges real-time NAS introduces to the practice of composers and performers.

4.2.2 Composing with Real-time NAS

As discussed in Section 2.2.1, RAVE learns to reconstruct sounds by compressing them into a continuous latent space, a multidimensional representation of learnt sonic features located in between an encoding function compressing the audio stream, and a decoding function reconstructing the sound. RAVE models can be interacted with in two ways: first, by using a sound source as input, with the model transferring the timbral features of the dataset into it, or, second, by bypassing the encoding function and forwarding control signals directly into its latent dimensions. These two methods may also be combined, for instance, by using one’s voice as input and adding control signals to offset the latent input values.

While the direct manipulation of the model’s latent space opens up significant creative potential, it also foregrounds a series of compositional and performative challenges that uniquely shape artistic practice with NAS compared to other synthesis methods. Of these, three have stood out through collaborations with artists from our lab’s extended community, and in my own artistic practice:

Arbitrary Encoding: The mapping of specific sound features (e.g. pitch, timbre, dynamics) to latent dimensions is autonomously determined by the model during training, and is different with every training run and dataset. Once a model has been trained, the distribution of these features is understood by the user retrospectively through direct interaction and exploration. In this sense, one approaches the model similarly to how a practising musician learns a new instrument, and often adapts the compositional strategy to accommodate discovery and hands-on engagement with the system.

Feature Entanglement: Individual latent dimensions typically encode multiple, intertwined sound features. For instance, the first latent dimension of a given model might correlate with amplitude and brightness, while the second might affect both frequency and noise level. This entanglement prevents independent access to, or representation of, individual sound features. Consequently, rather than abstracting individual sound parameter changes, it is more practical to describe or prescribe latent value changes.

Dimensional Entanglement: Latent dimensions are highly correlated with one another. This means that by changing the state of one latent, the response of all the others changes. This second type of entanglement makes one-to-one mapping strategies inefficient. Furthermore, as discussed in Section 2.2.1, the latent space is usually high-dimensional (in the case of RAVE, typically four to thirty-two dimensions), making it impractical to control all latents individually and at the same time. For these reasons, a synthetic approach to mapping can prove more effective, in which multiple latent dimensions are controlled with unified gestures.

Arbitrary encoding, dimensional entanglement and feature entanglement position the interface at the heart of the compositional process with real-time NAS. Rather than resorting to notational systems abstracting individual sound parameters, or impractically attempting to define univocal mappings between control parameters and sound features, performers and composers often benefit from exploratory and embodied approaches, in which sonic outcomes develop as trajectories through the latent space.

Crucially, as discussed in Section 4.2.1, inherent scores introduce a highly open and exploratory character as well. This parallel between NAS and inherent scores' exploratory character prompted this chapter's investigation into how the latter might be leveraged to streamline performance and notation with real-time NAS. I addressed this challenge by developing a novel type of inherent score for latent navigation in NAS (Privato et al. 2023b), and along with it, the notational strategies described in Privato et al. (2024).

4.2.3 *Magnetic Scores*

To operationalise and investigate the notion of inherent score through artistic practice, Enrique Tomás develops a series of tangible scores, wooden boards featuring touch-sensitive, engraved

inscriptions that the performer plays through tangible interaction (Tomás and Kaltenbrunner 2014). While tangible scores hold a visual dimension in that they suggest specific gestures through graphic inscriptions embedded in the instrument, their focus lies in the relational and situated emergence of the inscription. In other words, in these types of notational systems, the interface functions as a juncture between the intentions of the composer and the performer through a tactile element.

The instruments presented in this section share the relational and emergent nature of the inscription that characterises tangible scores, but instead of encoding the information through engravings on a board, they primarily rely on the interactions between a series of magnetically inscribed objects.³² For instance, by placing magnets of different dimensions, shapes and polarities within a wooden board such as the one shown in Figure 4.2, the composer defines the scope of the performer's gestures emerging out of the embodied interaction with the magnetic fields. Whereas tangible scores explore the boundaries between instrument and score by defining a space of possible action through a combination of visual and tactile elements, magnetic scores achieve this through somatosensory feedback (Rincon-Gonzalez et al. 2011), which depends on the user's proprioception in relation to the forces of attraction and repulsion exerted by permanent magnets.

This proprioceptive quality differs from the tactile one of tangible scores in that, as Brian Massumi argues, it folds tactility onto itself and introduces a dimension of medium depth between epidermis and viscera in which muscles and ligaments “register as conditions of movement what the skin internalizes as qualities” (Massumi 2002, 58). Similarly to how the hardness of a floor is understood as a resistance that enables standing and movement, or the softness of a cat's fur favours the motion of the hand, the forces of attraction and repulsion between magnetically inscribed objects define certain potentials for musical action.

To capture this proprioceptive dimension and distinguish it from the tactile one of tangible scores, and the different ways in which they inherently inscribe the performative

32 Magnets are ubiquitous in audio electronics (e.g. pickups, speakers), but their use as structural elements in score and interaction design, especially for gestural control, remains relatively unexplored. Recent examples include: Magalhaes' *Mobile*, featuring magnets on piano strings for bouncing/glissando effects (2012); Lepri's *Chowndolo*, a magnetic pendulum controlled by permanent magnets (see Section 4.4); and Wintergatan's *Marble Machine* (2016), in which ferromagnetic marbles interact with surfaces, timing controlled by magnetic step sequencer. In NIME proceedings, magnets have been used as position tags for precise detection at short distances (Paradiso et al. 2001; Pardue and Paradiso 2002; Mannone et al. 2018). Magnetic scores differ from these approaches in that they embed passive magnetic components into the sensors themselves, enabling somatosensory interaction through magnetic forces (repulsion in *Thales*, attraction in *Stacco*) rather than using magnets solely as position tags or actuators.

information, I describe Thales and Stacco, the two interfaces I introduce hereafter, with the umbrella term *magnetic scores*.

4.3 Thales

Thales (Figure 4.1),³³ the first of the two NIMes I introduce along with their notation, is named after Thales of Miletus, who according to Aristotle described the lodestone as “having a soul” (Raven 1957). The philosopher inferred the presence of a soul in loadstones from their ability to produce movement and affect each other, a suggestive idea that this interface poetically builds upon.

This effect is achieved by creating the impression of an invisible object that the user instantiates in the act of performing, momentarily existing as the consequence of the interaction between the reciprocally repelling magnetic forces of two hand-held controllers. In a preliminary evaluation with seven participants, users experimenting with Thales have described this emergent invisible object as a “soft, bouncing ball,” a “slippery surface,” a “soft bag that jumps,” a “solid sphere,” and a “ball filled with water” (Privato et al. 2023a).



Figure 4.1. Thales

Thales’ notational counterpart consists of magnetically inscribed boards. These are surfaces in which I embed magnets of varying polarities, dimensions, and positions, whose fields are revealed somatosensorially when the controllers are brought close to the board. The scores

³³ A demonstrative video of Thales and Stacco is available at <https://zenodo.org/records/20497498>

typically display no instruction, and provide a space for exploration in which the forces of attraction and repulsion of the hidden magnets guide the hands. Section 4.7 returns to a specific instance of these boards to discuss the interface's figurative semiotic level (see Section 3.4.1).



Figure 4.2. *A Magnetic Score*

4.3.1 *Hands Holding the Void*

Many factors contribute to the inception and evolution of a musical interface. One gets inspiration from diverse, sometimes eerily converging aesthetics, practices, theories, coming from far in time and through unpredictable routes. Not to mention the materials encountered by choice and by chance, and that one moulds, recombines, and thinks through. Thales is, of course, no exception.

In Alberto Giacometti's 1934 sculpture *Hands Holding the Void* (*The Invisible Object*), a bronze, stylised female figure sits on a throne, its hands facing each other as if holding an invisible object (Figure 4.3). I like to imagine Marco Donnarumma, the Italian performance artist I met on the first day of my PhD, looking at this statue, and walking around the female figure at MoMA where it is exhibited. That day in Reykjavík, Marco performed a piece titled *Ominous* featuring the *Xth Sense*, a biophysical instrument capable of amplifying sounds from

one's heart, blood and muscles (Donnarumma 2011). Inspired by Giacometti's sculpture, in his performance Marco tensed tendons and tightened muscles of arms and hands, activating the *Xth Sense* while theatrically moulding an imaginary sphere.

Sitting among the audience, I began wondering whether that sphere could somehow become tangible to the performer, while maintaining its fleeting, emergent and liminal nature. To conjure this object-phenomenon, it felt natural to take inspiration from Barad's *agential realism*. Quoting Barad, "reality is not composed of things-in-themselves or things-behind-phenomena, but of "things"-in-phenomena" (Barad 2003). This relational take assumes that no separation between the subject and the object of knowledge may be identified in the first place: they rather discursively and performatively emerge. In this sense, the coming-into-being of Thales' interface-space might be seen as a momentary stabilisation of a phenomenon, an *agential cut* that temporarily produces the distinction between a "this" and a "that," and thus makes it apparent to the senses.




Figure 4.3. From left to right, Alberto Giacometti's *Hands Holding the Void*, Marco Donnarumma performing *Ominous*, *Thales*.

4.3.2 Design

The reflections inspiring the interaction design of Thales inevitably infiltrated the design of the sensors as well. Rather than preserving the conventional subject-object distinction by using a magnetometer to detect the strength and position of external magnets, each of Thales' sensors embeds magnets as well, thus sensing variation in the surrounding magnetic field and at the same time actively engaging with surrounding ferromagnetic objects.

Thales consists of two 3D printed magnetic controllers with a diameter of 10 cm and thickness of 2 cm. Each controller mounts a three-dimensional gyroscope and accelerometer, a three-dimensional magnetic sensor, one microcontroller and a battery. At the centre of the controllers, right above the magnetometer placed on the PLA body, is mounted a cylindrical, vertically moving riser, with a diameter of 3 cm. The riser contains a neodymium magnet,³⁴ which is held in position above the magnetometer by the performer's hand holding the controller.

The orientation of the magnets inside the controllers is identical, so that when the controllers are held in front of each other, the magnets they contain repel and push on each other (Figure 4.4). This reciprocal repulsion turns the risers, otherwise unresponsive, into active components, causing them to push against the performer's hand. By releasing and modulating the pressure, one can extrude the riser and distance the magnets from the magnetometer, thus activating and finely controlling the readings of the vertical axis in the two controllers while changing the somatosensory perception of the interface-space.³⁵ Drawing a parallel to one of the traditional musical instruments I have a history of practice with, the manipulation of Thales' riser shares similarities with how one uses a trombone's slide, albeit on a minute scale.

Through a dedicated Wi-Fi protocol,³⁶ a wireless network is instantiated between the microcontroller mounted on each of Thales' controllers, and a third one connected to the laptop's serial input and acting as a server. Each controller forwards to the laptop nine data points, three relative to absolute orientation, three to acceleration, and three relative to the magnetic field strength on each spatial axis.

34 Neodymium is a rare metal known for its magnetic properties, a better choice than ferrite for building strong permanent magnets.

35 The proximity of the magnets to the sensor (< 6 mm) saturates the reading. In Thales, this limitation becomes a feature: when the readings saturate, the DSP controlled by the riser is off; when the riser is distanced above the 6 mm threshold, it activates and modulates the DSP.

36 The ESPNow protocol implemented in Thales offers very low latency (by specification, a minimum of 2 milliseconds) and high stability. For details see <https://www.espressif.com/en/solutions/low-power-solutions/esp-now>




Figure 4.4. Thales' magnetic interactions. Each controller reacts to and reads the magnetic fields of the other.

4.3.3 Sound Design

Between the two interfaces described in this chapter, Thales is certainly the one that more closely maps to the definition of a composed instrument (see Section 4.2.1). Indeed, despite a few initial tests with different sound engines and NAS models, throughout my practice I have been consistently mapping the controllers to one particular RAVE model that leverages the instrument's magnetic materiality as a sort of compositional motif.

The dataset I curated for this model consists of about one hour of recordings from magnets of different shapes and dimensions, thrown, rubbed, rolled on wooden and metallic surfaces and on one another. This compositional and design strategy strengthens the conceptual bond between the synthesis engine and the material components of the instrument. Over the two years after Thales' inception, I iteratively built upon this close pairing, incorporating new features and functions in a manner similar to how a guitar player experiments with various strings, picks, and pedals, while the instrument remains substantially the same.

Each of Thales' controllers is mapped to one instance of the RAVE model I trained on magnet sounds, of which I included only the decoding function. Each one of the decoders exposes a total of eight latents (the linear dimensions that together make up the latent space)

that can be individually manipulated with control signals. RAVE organises latents progressively from the most to the least meaningful through principal component analysis (PCA). I thus mapped to the first two latents respectively the x and y spatial movement readings from the magnetometer, followed by the x , y , z readings from the gyroscope, and finally those from the x , y , z axes of linear acceleration. The remaining data point, relative to the magnetometer's vertical readings controlled via the movement of the riser, is mapped to a series of additional Digital Signal Processing (DSP) techniques: a pitch shifter for the left controller, and an FM synthesis modulator with the model's output as carrier for the right one.

RAVE models trained on datasets consisting of sounds with short transients, when used without an encoder, present a characteristic granulated quality due to the repetition of identical blocks of sound. This becomes all the more evident when the control signals coming from the interface are relatively steady. For this reason, I offset each of the latents with a series of Low Frequency Oscillators (LFOs). This introduces a slow, looped dynamic transient, a sort of deep, prolonged, earthly breath that represents one of the signature sounds of Thales. Finally, to effectively control the amplitude while minimising the latency from the model, I added an independent envelope after the decoder's output on each of the controllers, triggered and modulated by the variation in the readings of the magnetometers or the amount of movement being exerted. As a result, the absence of magnetic interaction produces silence, while the sound is elicited by rapidly moving the controllers toward each other or any other magnetically inscribed surface.

4.4 Stacco

Stacco (Figure 4.5),³⁷ the second magnetic score I introduce, embeds Thales' magnetic sensors underneath an engraved surface. It inherits the design features of two instruments based on permanent magnets: on the one hand, Thales, with which it shares the entanglements of sensors with permanent magnets; on the other hand, the Chowndolo,³⁸ a magnetic pendulum whose trajectories are affected by permanent magnets placed over a metal board (Merendino et al. 2023). Stacco similarly exploits the liveness of ferromagnetic objects to manipulate sound,

³⁷ A demonstrative video of Stacco and Thales is available at <https://zenodo.org/records/20497498>

³⁸ The Chowndolo is an instrument of Giacomo Lepri's design. Giacomo joined the Intelligent Instruments Lab at the end of my first year of PhD. Together, we designed and developed Stacco.

while introducing an original take on the notion of inherent and magnetic scores through a notational approach that in Section 4.4.3 I will describe as *embodied sketching*.

The performer interacts with Stacco by throwing, displacing and combining magnetic marbles of variable dimensions and shapes on a wooden, oval surface with four magnetic attractors mounted underneath. Whereas Thales was designed to metaphorically translate the entanglements of RAVE's latent dimensions into the tangible manifestation of the controllers' intra-action (see Section 4.3.1), the spheres moving within and around Stacco's magnetic fields control the latent dimensions directly, while influencing each other's movements and unpredictably structuring into larger shapes.



Figure 4.5. Stacco V3

4.4.1 Design

Stacco consists of an oval laser-cut case containing a Bela board (McPherson et al. 2016) and four sensors, each combining one three-axis magnetometer with a 30×5 mm diameter neodymium magnet in a 3D-printed enclosure. The attractors are placed in four symmetrical points below an engraved board mounted on the top of a CNC'd body, which features a raised oval edge. Each magnetometer performs three-dimensional readings of nearby magnetic fields,

for a total of twelve continuous parameters that reflect, as in *Thales*, the strength of nearby magnetic fields. Circular engravings on the upper face of the instrument denote the position of the four attractors.

Over the course of these three years of research, *Stacco* has undergone three major iterations. These were partly driven by the need for increased precision and dependability, and partly by the growing interest in the instrument among the artistic community, with customised versions commissioned to meet specific artistic and technical needs. The first working prototype features an entirely laser-cut body, with a living hinge structure and the bodies of the 3D-printed sensors screwed to the back of the board on top of which one performs. The second version offers a drastic upgrade from the first, with movable attractors suspended under the board, and a CNC wooden body. The third version, developed specifically for the Barbican Centre exhibition *Feel the Sound*, embeds in the body of the instrument a series of four sliders mapped to the amplitude of as many RAVE models (Figure 4.5).

4.4.2 Sound Design

As discussed in Section 4.3.3, the coupling between the interface and the model in *Thales* has remained fundamentally unaltered since its inception, and iteratively developed through the years. However, *Stacco*'s mappings have changed considerably in the past three years, reflecting the diverse contexts of use and artistic needs the interface responded to, many of which are presented here and in Chapter 7. Broadly speaking, in my practice, as well as throughout collaborations with artists interested in experimenting with *Stacco*, two approaches to sound design with RAVE have emerged.

The first approach involves pure latent navigation (see Section 2.2.1). Here, the readings from *Stacco*'s attractors, often excluding the vertical axis (and thus, a total of eight data points), are mapped one-to-one to the latents in the decoder, thus bypassing the encoding function as in *Thales*.³⁹ In a few models, the first latent, typically loosely correlated with amplitude, is mapped to the total amount of energy present in the system, a strategy that drastically changes the interaction and invites a much more pronounced gesturality.

³⁹ I am currently exploring the vertical axis using magnetic rings. This affords a theremin-like navigation, but on a smaller scale and with enhanced positional precision due to the tangible response of the magnetic fields.

The second approach features instead the combination of timbre transfer and latent navigation. Here, one typically plays an instrument, an audio file, or sings into the RAVE encoder, and offsets its output values by manipulating the latent dimensions via Stacco before the decoding phase. This method introduces higher unpredictability, since a certain sonic outcome depends on a more complex combination of factors, an aspect that drastically reduces the impression of control (see for instance Section 7.2.1).

4.4.3 Embodied Sketching

As with all inherent scores, Stacco builds on the coupling of the interface with the notation. However, such coupling is not as tight as in other inherent scores, in which changing the notational layer typically implies a redefinition of the interface. Stacco’s magnets do not need to touch the instrument to influence the readings of the attractors: leveraging this action-at-a-distance, along with Giacomo Lepri I designed a series of oval cardboard and tracing paper sheets to be placed and drawn upon Stacco’s surface while playing. Composers and performers use these *embodied sketches* (Figure 4.6) as a notational technique to describe, prescribe, recall, or navigate meaningful areas in the latent space (Privato et al. 2024).

As discussed in Section 4.2.2, one of the consequences of NAS’s arbitrary encoding is that composers learn how the model represents sounds in the latent space through direct engagement. Thanks to the compositional affordances of embodied sketches, relevant areas in the latents can be at once discovered, incorporated into the score, and recalled, and scores can be dynamically replaced, overlapped, and reworked into articulated pieces both on and away from the instrument depending on the requirements of the model in use.

Embodied sketches also align with the feature entanglement constraints, which hinder the notation of individual sound parameters in favour of gestural and graphic notational methods. Finally, the model’s dimensional entanglement finds a parallel in Stacco’s ergonomics (see Section 2.3.1), since each attractor on Stacco can read three spatial parameters, allowing the user to control three latent dimensions with one single magnet. This synthetic approach to latent manipulation further extends through the entanglement of the different magnetic components, with the magnetic spheres influencing each other through proximity when placed on the attractors. Importantly, as emerging from workshops (Lepri et al. 2024), this entanglement tends to be reflected in the notation of the sketches, with fewer signs defining strategies for a large number of latents (Privato et al. 2024).




Figure 4.6. Embodied Sketch on Stacco, in Privato et al. (2024).

4.5 Study 1 - Thales: XAI and Sound in Space

Having introduced the two interfaces I designed around NAS along with the compositional strategies I developed around them, the second part of the chapter focuses on their application to investigate XAI questions. The aim here is to show how the outcomes from the studies informed the shift of the investigation toward what in Section 2.3.5 I termed Un-XAI.

The first study, featuring Thales, focuses on sound spatialisation, and explores through artistic means whether and how this dimension might contribute to one's understanding and navigation of latent distributions in autoencoders (Armitage et al. 2023). As in the case of much of the research presented in this dissertation, this investigation took shape out of a sustained collaboration with an artist interested in designing and performing a new instrument prototype.

4.5.1 Research Ideation

The artist I collaborated with on this occasion, Majella Clarke, expressed interest in augmenting the baton she uses to conduct classical orchestras. Using the Lab’s *Organium*, a library of technical elements for thinking-through-design (Magnusson et al. 2024), we augmented three batons with an accelerometer and microcontroller wirelessly connected to a computer running a series of RAVE models (Clarke 2025).

We used the *sonic batons* in a performative experiment, all three wielded to navigate the latent space of as many RAVE models’ autoencoders, one model for each baton. As in Thales, a RAVE decoder was fed with control signals derived from the accelerometer and gyroscope mounted on the baton, for a total of six data points per baton. No signal was forwarded to the remaining dimensions (two on each of the models we used).

The performance involved Majella, myself, and my colleague Victor Shepardson. In conversations afterwards, Victor and I observed how, by pointing the batons at precise locations in the physical space, we were consistently setting spatial anchors to particular latent configurations and sounds. The spatial dimension seemed to facilitate our intuitive navigation of the model’s high-dimensionality, enabling the recurrence of certain sounds and the repetition of patterns, which in turn facilitated meaningful musical dialogue. However, the setup provided by the venue was stereophonic, which complicated the aural localisation of the sounds. This led us to reflect on whether a three-dimensional projection of sound might facilitate the understanding and exploration of the model’s latent distribution. To investigate this, I needed access to an immersive sonic environment, ideally a spherical one with a high number of speakers, that would afford a very precise localisation of the sound sources.

4.5.2 DodekaOTTO

I was granted access to an immersive sonic environment on the occasion of an Erasmus+ research residency at the Tangible Music Lab, in Linz. The Lab’s system, named DodekaOTTO (Figure 4.7), is a sound sculpture for spatialised music consisting of 20+1 3D-printed speakers and a lightweight aluminium structure in the form of a dodecahedron (Mitterhuber et al. 2025).

I implemented a real-time spatialisation algorithm using ICST Ambisonics (Schacher 2010), a series of MaxMSP externals for 3D sound processing based on the Ambisonics open-source protocol (Zotter and Frank 2019). Due to the high dimensionality of the latents per

single model, I decided to experiment with Thales' controllers rather than the batons, since these allow controlling a much higher number of dimensions. Throughout the residency, I experimented with the magnets RAVE model and with a choir model, the latter provided by artist and colleague Jonathan Reus (Reus 2023).⁴⁰ For the magnets model, I used the mappings described in Section 4.3.3, with one model per controller. The latents exposed by the choir model (eighteen, matching the number of data points controlled by Thales) were instead mapped one-to-one to the decoder latents.

Throughout the residency, I experimented with a few different mapping configurations across both models, but consistently mapped the Ambisonics space to the gyroscope values in the controllers, which in turn affected the first three latents in the models. With the two magnets models, each controller was mapped to one sound source in the Ambisonics space; with the choir model, only the left controller's spatial readings were mapped to a sound source. Rather than tweaking my algorithms throughout the week of sessions, I chose to focus on practising with them. This was crucial for the insights discussed in the next section.⁴¹



Figure 4.7. Performing with Thales and DodekaOTTO at Ars Electronica.

⁴⁰ The model was trained by Jonathan Reus on a collection of open choral datasets augmenting core recordings made in collaboration with the Musilon student choir of the University of Twente.

⁴¹ Videos from the practice sessions are available at <https://zenodo.org/records/20497433>

4.5.3 *Insights*

During the week of practice, it became clear that learning to navigate RAVE with Thales through the DodekaOTTO system was quite different from my experience with the baton in a stereophonic environment. Furthermore, while the spatial projection of the sound seemed to offer a slightly clearer anchoring to certain trajectories and sonic gestures, consciously mapping the high-dimensional algorithmic spaces onto the three-dimensional ones remained challenging. If the immersive system had contributed to better understanding the model's latent space, it had done so through sustained practice and on an intuitive level.

Consider the high number of dimensions we navigate as we speak: tongue, lips, teeth and jaw position, diaphragm tension and flow modulation, relaxation of the throat and nasal cavity, to name a few. As long as we do not explicitly “think” of this n -dimensional space, as long as we do not attempt to articulate it in the act, we can speak fluently. Similarly, navigating the high dimensionality of the NAS models became much simpler once I gave up on trying to consciously map every dimension. It was only when I shifted from embodied processes into analytical ones, when I distanced myself from the intrinsic high-dimensionality of experience and the inherent “messiness” of being in the world, that I experienced the model's high dimensionality as a problem to solve.

While the extra spatial dimension offered by the DodekaOTTO contributed to lowering the steepness of the learning curve, what made the most difference was ultimately sustained practice. Throughout the regular practice sessions that accompanied my week of residency, I began understanding RAVE, Thales and the DodekaOTTO as a large-scale musical instrument I was controlling from the inside. Using Thales' magnetic interactions, I intuitively learnt to negotiate gestures and sounds between the inner workings of the model and the outer workings of the DodekaOTTO. As the practice progressed, these negotiations found a synthesis in an embodied, latent topology that I could hardly describe in plain language, but could intuitively navigate.

Another relevant and unexpected aspect that influenced the learning derives instead from the idiosyncratic semiotics of the interface I had designed, and, specifically, what Scolari describes as the plastic level (see Section 3.4.1), regarding how the material features of the interface (e.g. shape, colour, material) contribute to processes of sense-making (Scolari 2009). In comparing the use of Thales and my previous experience with the sonic baton, I realised how their different plastic features were having a critical impact on my playing, with their

shapes informing an entirely different thinking of sound and gesture, and ultimately of how I thought of the latent space of the RAVE models.

The shape of the baton resembles a line or an arrow. It is designed to point at sections of musicians and bodies in the orchestra, and to convey the intent of the conductor as accurately as possible. It may be described, as it were, as an “outward-looking” device. On the other hand, the circular shape of Thales balances centripetal and centrifugal forces. The circular *mandala*, for instance, is a highly productive metaphor for Pauline Oliveros, one that informs the relation between the inner experience of the self with the outer world, as well as that between focalised and open listening (Oliveros 1984). In using the baton, it felt natural to assign different sonic outcomes to the specific areas in the performative space I was pointing to. With Thales, the latent space became instead less of a literal metaphor and more of a proprioceptive one, with the latent dimensions of the model and the Ambisonics space collapsing from opposed directions (outside-in and inside-out) towards my hands.

These two insights, namely that a high-dimensional latent space can be understood through sustained engagement without being formalised or conceptually reduced (perhaps a low-hanging fruit for practising musicians, yet an underdeveloped aspect in XAI research, which as discussed in Section 2.3 tends to favour top-down approaches), and the capacity of the interfaces I performed with to translate in different ways the workings of NAS models, lay the foundations for a second study using Stacco, and for the discussion I offer in Section 4.7. In the study, which focused on embodiment and AI explainability in artistic practice via embodied sketching, I looked into (and partly challenged) the common XAI assumption that a regularised latent space tends to be preferred to an unregularised one.

4.6 Study 2 - Stacco: XAI and Latent Regularisation

The second study I introduce in this chapter focuses on the embodied perception of latent dimensions in NAS. The study was conducted together with my colleague Victor Shepardson, who is responsible for the development of the technical approach for latent space adaptation described in the study design (Section 4.6.1). Together, we designed the study methodology and analysed the data, and with Giacomo Lepri presented the outcomes in a NIME paper (Privato et al. 2024).

The aim of the study was to investigate, through the embodied and pragmatic perspective described in Section 2.3.4, (i) whether performers can recognise a consistently regularised latent space throughout different models played on the same DMI, and (ii) whether and how such latent regularisation benefits their expressive and artistic needs.

To answer these questions, Victor Shepardson developed an adapter module enabling the transferability of the latent mappings from one model to another, and thus the predictability of the latent distribution of model B having learnt model A. This approach follows a common line of research in XAI literature, where a linear transformation layer (what is here termed *adapter*) is introduced between the latent representations of two different models, enabling the encoder of one model to communicate effectively with the decoder of another while preserving the spatial relationships through different latents. In other words, performing the same gesture from one model to another when the latter is adapted to the former results in relatively similar sonic outcomes.⁴² We embedded the adapter in Stacco’s algorithm, and designed a study involving embodied sketching to facilitate notation. In the study, we invited artists to compose and repeat a series of musical pieces on Stacco using randomly adapted and unadapted models.⁴³

4.6.1 Study Design

We recruited five participants, P1-P5, from the artistic community of the area. Four of them are trained musicians, each with a different musical expertise: P1 is a trained classical pianist; P2 is a bass player; P3 and P5 are electronic musicians. Of the five, only one (P4) is not a trained musician, being instead an installation artist. We planned five individual sessions, preparing Stacco with a blank oval cardboard for embodied sketching, a pencil and a rubber on the side to notate the score, and four magnetic spheres each placed at the centre of one attractor.

We began each session by introducing Stacco as a new musical interface of our design, which we are experimenting with to investigate a novel synthesis technique called “neural synthesis,” and invited the participant to explore and play the instrument. In this phase, the

⁴² I write “relatively” because differences in the data when adapting to each other models trained on very different material (say, drones on the one hand and percussion on the other) also influence the perception of similarity (see Section 4.6.5).

⁴³ For a technical description of the adapter developed by Shepardson, refer to the conference paper (Privato et al. 2024).

participant was free to ask questions about the instrument’s workings. At the end of this exploratory phase, we invited each participant to work for 10 minutes on one or more small pieces, indicating a total length of two minutes. We also clarified that the pieces should be repeatable, and that to help memory one could choose to notate the piece on the available embodied sketch.

At the end of the ten minutes, the participant was invited to perform the composition three times, first on setting A, that is, the same model used to compose, then on two different settings (B and C). After each performance, we asked the participant to freely comment, and after playing with B and C we asked what, if anything, from the original composition was preserved in playing on the new settings.

Finally, we asked two more questions: “Which of the two settings (B and C) preserved most of the original character of the performance?” “Which setting did you enjoy most, and why?” The whole process was repeated once more in a second round for each participant, with a different series of settings A, B and C.⁴⁴

4.6.2 *Settings*

In each of the ten total rounds, setting A consisted of one RAVE model, while settings B and C both used a second, different model, with either B or C including the adapter module that made the latent space more similar to that of A. With this setup, one would find a gesture performed in A to produce similar sonic outcomes when performed in the adapted setting rather than the unadapted one. The order in which the adapted model was presented (as B or C) was randomised to reduce the influence of order effects. Each participant encountered settings derived from RAVE models trained on the same data in one round, and different data (between A and B-C) in the other round. The order of these same and different rounds also varied between participants.

Setting A always used a RAVE model trained on guitar or saxophone sounds. When B and C used models trained on different data, it was the organ one. By using three random seeded versions of each guitar and sax model, we were able to use a unique configuration of A, B, and

⁴⁴ I write P3-2, for example, to denote performer 3, round 2.

C in every round. This made our observations less dependent on the particularities of any pair of RAVE models (see Table 4.1).

	Setting A	Setting B	Setting C
P1 - 1	G0	G1	G0 → G1
P1 - 2	S0	S0→OR	OR
P2 - 1	S1	S1→S2	S2
P2 - 2	G2	G2→OR	OR
P3 - 1	S2	S2→OR	OR
P3 - 2	G1	G1→G0	G0
P4 - 1	G2	OR	G2→OR
P4 - 2	S0	S1	S0→S1
P5 - 1	S2	S0	S2→S0
P5 - 2	G0	OR	G0→OR

Table 4.1. Study Settings, rounds 1 and 2. P1 to P5 rounds 1 and 2, with guitar (G), saxophone (S), organ (OR) trained on random seeds 0 to 2 and latent adaptation.

4.6.3 Results

The effect of the adapters in transferring spatial relationships from one model to another was in most cases evident. This was for instance the case with P3-2, who described the adapted setting (B) as fun and “more trustworthy” than the unadapted one, even responding with a “neater” feeling than setting A in navigating the latent space. Similarly, P5-2 described the unadapted setting (B) as very different from A, and the adapted one as a “blend between A and B.” In most cases, when asked which setting they enjoyed the most, the performers described the adapted model as more fun than the unadapted one. This suggests that even when entering embodied experience through an interface such as Stacco, the adaptation of a model’s latent space to another is indeed generally perceptible.

However, in a smaller subset of cases, participants perceived no particular similarity between setting A and the adapted setting B or C. P1 even rejected any expressive or tonal similarities of settings B and C with setting A, arguing that what had remained consistent was

the dynamic response of the spheres to the gestures (the overall energy in the system that we mapped to the first latent by study design). This was also the case with P5-1, where settings B and C were perceived as similar to each other and quite different from setting A. Furthermore, P3-1 perceived the unadapted setting (C) as more similar to setting A in that it provided higher responsiveness to their gestures, thus making the experience more enjoyable. At the same time, the adapted setting (B) was described as disappointing in that it compelled P3 to change the gestures. These observations suggest that once it enters embodied experience through an interface such as Stacco, the algorithmic adaptation of the latent space of different models can be less easy to perceive.

4.6.4 P2 and P4

Interestingly enough, when asked which setting maintained most of the original character from their piece, P2 and P4 reported the unadapted settings to better preserve their composition at least once. For this reason, it is worth detailing here their compositional approaches and responses.

P2 engaged with the composition task by developing an idiosyncratic vocabulary of gestures, and performed in a semi-improvisational manner, continuing to experiment long past the suggested two minutes. P2 was also the only participant who did not notate the embodied sketch with the pencil, preferring to rely on muscle memory.

The first round was particularly noteworthy, since the similarity in sound between A and B (which had been adapted) was evident both to myself and my colleague Victor, and, crucially, was also commented on by P2 themselves. However, P2 identified the unadapted C as “paradoxically” better at preserving their original composition: in the adapted setting (B), the transients were perceptibly preserved along with the overall character of the composition, but the unadapted setting (C) provided a darker, “sombre change of mood.” Despite the tonal similarity of the adapted model and the repeatability it clearly afforded, with the unadapted one P2 felt more capable of going “where [the] inner ear was leading to.” Similarly to the first round, in P2-2 the participant struggled to reproduce the original compositional idea with the adapted model, but enjoyed distancing from it and exploring the uncharted territory offered by the unadapted one.

P4 resorted instead to an extensive use of the embodied sketches, drawing an articulated and fine score while systematically exploring the attractors in sequence. In both rounds, P4 chose the unadapted model as the most effective in preserving the character of the piece, and even gave specific reasons as to why it was more similar to setting A, describing specific gestures and how the sound they produced had been preserved.

The participant described the unadapted setting as more sensitive and therefore fun in round 1, and as more similar in character to A, in that the sounds it produced were “longer” than A. Interestingly, P4 added that they improvised more on B and had more fun because the sounds “were more out of control.” P4’s experience illustrates how, as we interact with a model through a DMI, the frictions from gesture to latent space can be more relevant and interesting to the user than the similarities in the mappings between different sound engines.

4.6.5 Study Conclusions

Due to the limited number of participants, findings from this study should be seen as preliminary, and warrant further investigation with larger cohorts and different use cases. Nonetheless, two conclusions in response to the initial questions can be drawn. First, once it enters embodied experience through a NIME such as Stacco, the algorithmic adaptation of latent space can become less easy to perceive. This does not mean that such methods, or any other XAI implementation applied to latent distributions, should be discarded altogether. The possibility of transferring one’s embodied knowledge from one model to the other can be appealing in many performative contexts. On this note, it is nonetheless relevant to stress that adaptation techniques such as the one implemented in this study usually apply to models trained on different datasets, which from the outset might exhibit limited and sometimes not so apparent similarities. Although the similarity between the mappings of the models used in the study was verified, outside of the experimental setup this can represent a bottleneck for transferability.

Second, and most importantly for the purpose of this dissertation, an adapted latent space might not necessarily be valued by the composer, who might feel more attuned to a different yet more sensitive setting, to a richer sound over pattern repeatability, and value surprise and instability over control and predictability. P2, in particular, explicitly framed openness and emergence as an inherent and necessary element for their piece, a consideration that clearly operationalises Eco’s openness of the second degree (see Section 2.4).

4.7 From XAI to Un-XAI

In this final section, I discuss the insights gained from the studies on XAI, and in a trajectory that Chapter 5 further consolidates, I begin the dissertation's overarching transition into the theme of spectrality. While Chapter 5 will demonstrate how this theme emerged in my own artistic practice but also broadly applies to the works of other artists focusing on music AI, here I outline how my research focus began transitioning from XAI to Un-XAI, and to the epistemological relevance of what resists explanation. To ground the discussion, let us begin by outlining the three insights emerging from the studies:

First, the different shapes of the interfaces I performed with (the circular shape of Thales, which I compared with the arrow-like shape of the conductor's baton) informed different understandings of the model and of the latent space (see Section 4.5.2). When we investigate XAI questions, it thus becomes crucial to consider the interfaces that afford the interaction as part of the context of use, rather than observing the architecture in itself. The Stacco study in Section 4.6 operationalised this insight.

The fruitful analysis of the interface's plastic semiotic level (see Section 3.4.1) offered in Section 4.5 motivates and paves the way for the analysis of the figurative level in Section 4.7, and the communicative and meta-communicative levels of the interface in use discussed in Chapter 5.

Second, due to the high dimensionality of VAEs, the positional mapping of the latent space into the sound space is non-trivial. However, through sustained practice, one can develop and learn to navigate an embodied and tacit map that does not require consciously reducing the high dimensionality of the latent space into three spatial dimensions. As Michael Polanyi argues, "we know more than we can tell" (1966, 4). In other words, knowing is inherently multidimensional, and it is possible and even likely for us to know something on one level but not on another. To return to Derrida's critique of the metaphysics of presence (see Section 2.5.1), knowledge is rarely, if ever, fully present to us.

Third, XAI research usually considers a regularised latent space highly desirable compared to a non-regularised one. However, this is not necessarily true in artistic practice, since artists often value unpredictability, openness, and emergence as relevant to meaning-making and thus

integral to their compositions. These become crucial aspects to consider in discussing and investigating explainability questions.

Let us further elaborate on the latter point. It is relevant to note that the etymological root of the word “to explain” presents us with yet another spatial metaphor. *Ex-planare*, to flatten out, to make into a plane (a single, two-dimensional surface versus the high-dimensional spaces in which VAEs distribute information) implies a reduction in the dimensionality of a system. A two-dimensional flattening that makes it more understandable. Yet, as Eco shows, it is precisely from a system’s complexity and unpredictability, from the noise that infiltrates it, that unforeseen interpretations are allowed to emerge (Eco 1989).

That there exists an active interest among composers in systems that are complex to the point that “one can’t know what will happen” (Waters 2007, 8) does not entail that we should not attempt to explain the opaque workings of generative models. Rather, this insight points to the possibility that a space of not-quite-knowing (knowing at some level as in the Thales study, knowing in “this” instance but not in “that” as in the Stacco one) can indeed open up new potentials for artistic expression, and thus should not necessarily be resolved *a priori*.

As discussed in Chapter 2, Derrida’s hauntological corpus may be viewed as an overarching reflection on the necessity and inevitability of the liminal. His critique of the binary of presence and absence, both of knowing and being, in Western thought, points indeed to the inherently deferring nature of language and of ontology itself. In *Specters of Marx*, Derrida uses the term *exorc-analysis*,⁴⁵ to describe the attempt to explain in order to explain away, and even a certain hostility towards “the spectrality of the spectre,” its occupying a place beyond binary oppositions (Derrida 1994a, 58).

The point being, for Derrida, that living with the liminal (the not fully known and thus not fully present) is inevitable, and that the very evitability of this truth is a “staged” one. Rather than resolving the unresolvable by “exorc-analysing the spectrality of the spectre” (Derrida 1994a, 58), Derrida argues that it is crucial for the scholar of the future to learn how to let ghosts speak, or “how to give them back speech,” without making conversation with them, which would annihilate the ontological and epistemological distance the ghost embodies (Derrida 1994a, 221).

⁴⁵ In (1994b, 37), Derrida uses the term *exorcizanalyze*.

In my work with Thales and in the user studies with Stacco, what proved most compelling and productive was indeed this fundamental resistance to explanation. With Thales, I could intuitively navigate the model's implicit topology but could not conceptually grasp it. In the Stacco study, P2 and P4 evidently enjoyed not knowing each model at the outset, and valued the uncertainty as an aspect integral to their compositions. This compelled a reframing of my initial research on XAI, from a positive stance (assuming that, whether from a techno-centred or pragmatic angle, one should indeed exorcise-analyse the model), into looking at the capacity of what eludes full explanation (importantly, not at the absence of explanation) to produce new artistic practices and discursive metaphors. Chapter 5 will discuss some of these as they emerged in my practice, and from there observe the broader landscape of music AI.

4.7.1 Toward Spectrality

Transitioning from not-knowing into spectrality as per the dissertation's overarching journey implies expanding the discussion from the epistemologically to the ontologically liminal (see Section 2.5). While Chapter 5 will focus on how the latter entered my artistic practice, it is nonetheless relevant to show how a certain discourse around this theme emerged from the outset of the research, especially as I adopted principles from new materialisms for the design of the NIMes supporting my practice (see Section 4.3.1).

While it is beyond the scope of this dissertation to offer a comparative analysis of Barad's agential realism and Derrida's hauntology, it is nevertheless useful here to begin by reflecting on how these thinkers view what *is* from a relational perspective: for Barad, through the principle of intra-action, which rejects the existence of pre-existing entities, and for Derrida through the deferring nature of meaning and the rejection of transcendental signifiers. From different angles, these philosophies frame absence as a pre-condition of presence. Whereas Barad locates presence-absence in the relational and performative coming-into-being of matter, Derrida locates it in the inherent and performative deferral of signification.

My practice with Thales functioned in itself as a connecting thread between these two thinkers, with the things-phenomena that the controllers conjured by design persisting as the trace of a presence. Users experimenting with Thales have described the repulsive forces between the controllers as a "tennis ball," a "trampoline," a "soft bag filled with water" (Privato et al. 2023a). From a Baradian perspective, such tennis ball, trampoline and bag filled with water may be viewed as things-phenomena that emerge from dynamics of intra-action. But as

I iteratively and performatively constituted these things-phenomena in my practice, I inevitably conjured the eerie vitality of a movement with no agent, the trace of an absence, the fleeting “soul” described by Thales of Miletus.

This transition from the intra-active to the ontologically liminal became progressively clear in my practice, and as I reflected on the figurative semiotic level of the interfaces (see Section 3.4.1). For instance, as I worked on a magnetic score to perform with Thales, I ended up performing an Icelandic spell (Figure 4.8).

The score features a magnetic board embedded with magnets of different dimensions and polarities. The magnets are hidden underneath the board, but their positions and strength are revealed by a series of circular engravings connected with lines, as vertices of a triangle, a symbol I found in an Icelandic book of spells during the first months of my research at the Intelligent Instruments Lab (Skuggi 2015).⁴⁶ The proximity of the controllers to the circular engraving “activates the spell” during the performance, whose “effect” and “power” are perceived as the magnetic repulsion of the board and the controllers. This artifice turns the interaction design of the system into a metaphorical conversation with the otherworldly.

⁴⁶ The book prescribes engraving the spell onto a wooden board and placing the board under one’s pillow to favour sleep. An appropriate spell for a PhD candidate.



Figure 4.8. Thales' Icelandic spell score.

A similar tendency emerged with Stacco as well. I remember testing Stacco for the first time at the Intelligent Instruments Lab's workshop space together with Giacomo Lepri. It was early in the morning when my colleague Jack Armitage joined us as we hurled magnetic spheres across its surface. The marbles, launched parallel to the rim, refused to trace their natural trajectories, collapsing toward the invisible attractors like satellites in a decaying orbit, and accelerating as they spiralled around the centres. As more marbles gathered on the board, the complexity increased, with the magnets escaping their paths to collapse into and veer away from one another, as if moved by invisible hands. Jack was visibly amused by these interactions, at least as much as we were. In describing the marbles' counterintuitive movements (the same uncanny and seemingly autonomous agency which led Thales of Miletus to think of the loadstone as having a soul), he likened the instrument to a Ouija board, the well-known device for the communication with spirits.

These metaphors proved surprisingly persistent. Over the years, they reappeared in multiple and independent instances, gradually transferring from the instrument's plastic and figurative levels to the communicative and meta-communicative ones of the interfaces in use. While the latter two will be further expanded upon in the next chapter by discussing the

performances I ideated around my interfaces and the audiences' reception of them, it is nonetheless worth providing here one example stemming from a collaboration with Robert Laidlow, the artist who besides me has spent the most time performing and composing with Stacco.⁴⁷ Robert contacted me in the spring of 2025 to collaborate on a new piece. The piece, as he described it without knowing the direction that my research had taken, features the instrument used as a communication device, a Ouija board to all intents and purposes, through which one would interact with the spectres of their past selves.

In Chapters 1 and 2, we saw that spectrality as discussed here does not necessarily involve spirits, ghosts, spectres, or supernatural events of different kinds. Yet, it can and often does include such connotations as discursive metaphors. The lodestar, it is worth stressing, is here the definition provided in Section 2.5, and in particular the ontological and epistemological connotations of presence, which the visor and auditor effects mediate so effectively by recognising a gaze and thus a presence arising from that which is not known. It is through this linkage between liminal presence and knowledge that it becomes possible to look at Un-XAI, at Thales' tangible yet invisible object, at Stacco's Ouija-like shape, as branches stemming from the same tree.

4.8 Conclusions

This chapter contributed to the expansion of the notion of inherent scores, open systems in which the musical interface and the notation merge with fluidity, and to its application to NAS. The compositional and performative challenges of NAS identified in Section 4.2.2, namely arbitrary encoding, feature entanglement and dimensional entanglement, tend to favour notational methods grounded on exploration. This exploratory character is also a feature of inherent scores, where the inscription is interpreted through embodied engagement. I thus introduced two novel inherent scores, Thales and Stacco, which leverage this common exploratory nature and that I specifically designed to perform and compose with NAS. I framed these as magnetic scores, an instance of inherent scores where the inscription emerges relationally through the interactions between magnetically inscribed interfaces.

One performs with Thales by manipulating a relationally emergent “object,” resulting from the reciprocal intra-action between the repelling magnetic fields of the controllers, or

47 For a detailed discussion of how Robert used Stacco in his practice, see Section 7.2.2.

between the controllers and a magnetically inscribed surface. Stacco, featuring a board with four magnetic attractors for latent navigation through the displacement and combination of magnetic marbles, further builds on Thales' intuition. Whereas Thales features a highly composed nature, Stacco offers a higher notational flexibility through the notation of latent trajectories on the body of the instrument (embodied sketching).

I used Thales for a practice-based investigation of the spatial projection of NAS latent space in a multichannel environment. The investigation offered two contributions relevant to XAI. First, practice may be viewed as an explanatory process in itself, yet one in which explanations are not necessarily consciously formalised. Second, the materiality of the interface represents a highly productive semiotic site, that enters into a discursive relation with the model it controls, and should thus be considered when implementing XAI strategies.

Building on these insights, I used Stacco in a study on the embodied perception of adapted and unadapted latent spaces. The relevance of the study lies in two emerging insights, namely (i) that as it enters embodied experience, the regularisation of a VAE latent space can be less easy to perceive, and (ii) that even when the adaptation is perceived and apparent, participants may prefer an unadapted space, which preserves the unpredictability of the model's unmapped latent space.

Derrida invites the scholar of the future to learn to let ghosts talk, rather than conjuring them away. Following this invitation, as the core contribution of this chapter I argue for a reorientation of XAI priorities, one that includes the creative and epistemological potential of what resists full explanation.

While this chapter offered a few preliminary insights into the first research question, the next one will fully answer it by expanding into my artistic practice with Thales, Stacco, and distinct types of NAS models. This will also establish the necessary groundwork to address the second question, which concerns the development of a theoretical framework around these emerging themes.

5 Works in Movement

*The more enlightened our houses are,
the more their walls ooze ghosts.
Italo Calvino (1987).*

5.1 Introduction

In a residency at the Milanese Studio di Fonologia Musicale in 1957, Henri Pousseur worked on a piece titled *Scambi*, the Italian word for “exchanges” (Boulez 2001). To be activated, the piece requires rearranging a series of tape-recorded sonic layers. This makes *Scambi* a quintessential operationalisation of what Eco describes as *openness of the second degree* (see Section 2.4), in which aesthetic appreciation lies less in the final recognition of a form, and more in the apprehension of a continuously open process of discovery that allows for the contextual germination of poetic meanings (Eco 1989, 74). In *Opera Aperta*, Eco exemplifies one creative operationalisation of open works of the second degree through the notion of *works in movement*, a type of open work consisting of unplanned or physically incomplete structural units, reorganising themselves or to be reorganised in the context of the aesthetic enjoyment (Eco 1989).

Pousseur’s example of collaborative meaning-making through material rearrangements provides a fitting introduction to the three pieces I discuss here. I describe these under the umbrella phrase *works in movement* in that they incorporate various types of contextual rearrangements whose efficient cause extends beyond a composer-performer linear relationship. This includes the viewers’ direct or indirect activation of the space (see Section 5.3), but also extends to the physical and metaphorical movements exerted by non-human actors. The latter is of particular interest for this dissertation, and encompasses the active organisation of musical material by generative architectures (see Section 5.4), the seemingly autonomous material rearrangements performed by my interfaces (see Section 5.2), and even metaphorical “movements,” such as the discursive coming and going of the “souls” identified by Thales of Miletus (see Section 5.2.4).

In what follows, I introduce three of the works in movement I implemented as part of my artistic research. These pieces offer a series of primary and secondary contributions in relation to the broader aims of this dissertation. The primary contributions are a cluster of diverse yet intertwined perspectives on AI and spectrality, which further expand the insights offered in the concluding sections of Chapter 4. The secondary contributions consist of a series of reflections that are tangentially related to the dissertation's overarching narrative, but nevertheless relevant to music AI in general. These will be discussed in the corresponding sections.

In the first of the three works, titled *Mouja* (Section 5.2), through séance aesthetics the NAS model's weird algorithmic traces (e.g. choirs with impossible lungs, unintelligible voices, and digital artefacts found in the latent space) find an embodiment-without-presence in the invisible forces activating Stacco's marbles and the action-at-a-distance of Thales' controllers. From the plastic and figurative levels of the interfaces (see Section 4.7), the reflection on spectrality thus infiltrates the communicative and meta-communicative ones of their performative use (see Section 5.4.2).

The second work is a site-specific piece titled *End of Empire* (Section 5.3). This large-scale installation, featuring four interactive textile columns embedding NAS models, explores the use of generative technologies as part of community-oriented participatory processes. Here, I describe the process leading to the collection of the data as a collective, situated and participatory engagement with memory and belonging. After reporting on the experience of working on *End of Empire*, I reflect on how sonic fragments collected in participatory processes as mediated and rearranged by the NAS models can, given certain conditions, produce coexisting effects of familiarity and estrangement.

The third work, titled *Mouja+*, investigates AI, textuality and presence. *Mouja+* is a performance piece and a performative pipeline combining small-scale NAS models (RAVE) with a text-to-audio model (Stable Audio Open), whose complementarity (the former functions in real time and focuses on timbral control, the latter functions asynchronously and affords formal control) I outlined in Section 2.2. By adapting and combining a series of scores from Fluxus artists which, much like in Pousseur's *Scambi*, I live-prompt and layer onstage through my magnetic scores to develop a narrative, I investigate and reflect on two aspects: first, how the use of natural language changes as we engage with text-to-audio models; second, how a text-to-audio model can be used to playfully conjure onstage the spectral referents of bodies and actions that never were.

These three works demonstrate different modes of what in the latter part of the chapter I term *haunted work*, an extension of Eco's framework that increasingly characterises AI-mediated artistic practices, in which algorithmic indeterminacy grounds upon indeterminate and deferred presence. While the open work introduces different types of ambiguity to enable the emergence of multiple interpretations, in the haunted work presence and absence become the foundational productive ambiguity. This reflection bridges the epistemologically liminal, which Chapter 4 focuses on, with the ontologically liminal, as per the working definition of spectrality offered in Section 2.5.

The chapter concludes by observing how the notion of haunted work applies beyond my own artistic practice, offering an overview of pieces from artists working in different ways with AI, and from there into digital culture. By showing how a certain spectral discourse emerges with particular force in AI-mediated cultural production, this concluding section fulfils two aims. First, it demonstrates the necessity for a vocabulary to investigate the spectralising effects of generative technologies. Second, it sets the empirical ground for the theoretical investigation in Chapter 6, which will attempt to provide such a vocabulary through hauntography.

5.2 Mouja

The relationship between electronic communication technology and intermundane communication has deep historical roots, dating at least back to the introduction of Morse's electric telegraph and spiritualism's coeval growth in popularity (Sconce 2005). Thanks to a certain alignment with the principles of the nascent electrical science, the technologically mediated communication with the beyond could claim a distinguished status over more "superstitious" forms of mysticism (Packham 2024).

From one angle, this higher status may be warranted. As Roach argues in his investigation of cultural exchange along the Atlantic rim through theatrical, musical and ritual performance, modernity itself might be understood as a new way of handling (and thinking about) the dead, who have been segregated from the spaces of the living by the rationalising project of the European Enlightenment (Roach 1996). But as Stanyek and Piekut argue in their reflections on the dead spaces of the recording studio, new production technologies transform

the segregation between the living and the dead into a continuum, with agencies intra-acting across distant physical and temporal planes (Stanyek and Piekut 2010).

Similarly, if capital may be viewed as dead labour (Marx 1887/1995) the technical abstraction of human labour into increasingly automated processes transforms the social relationship of production into a liminal phenomenon. Technology, and particularly sound technology, which acts as a disembodied agent (see Section 2.5.2), becomes an effective mediator between the collaborative efforts of the living and the dead.⁴⁸ As a seminal example, Stanyek and Piekut discuss how, back in 1932, Caruso's 1907 recording of *Vesti la Giubba*, from Leoncavallo's *I Pagliacci*, was stripped of the unsatisfactory orchestral accompaniment and posthumously complemented with a novel one (Stanyek and Piekut 2010).

Through the recording apparatus, time stretches and distends. But whereas from a hauntological angle such expanded and diverging temporalities haunt the present through their unresolved tensions, in the case of séance-inspired aesthetics, techno-mediated ghosts become harbingers of artistic invention. Building on Eco and on the chapter's title, it may even be said that what infiltrates the signal, here, what "moves" the works, is the dead itself.

5.2.1 Concept

Mouja, the first piece I examine (Figure 5.1), launches this chapter's investigation into spectrality through séance aesthetics, understood as a compositional approach focusing on the performative conjuring of spirits (Packham 2024). As discussed in Section 4.7.1, the interfaces I developed quite naturally led to narratives of otherworldly communication among lab members and collaborators (particularly, in the case of Jack Armitage and Robert Laidlow). Since in my practice I began investigating and interpreting NAS' affordances through these instruments, to performatively engage with AI from this angle felt quite natural.

The distance between the recording studio dubbings discussed above and music AI is indeed shorter than one might think. AI translates memory, and thus the deadness of the dead, into a searchable latent space of present potentials. Yet, the generative architecture reconfigures the sonic material on its own terms, through recombinatory dynamics based, rather than on causal relations, on machinic correlation processes that organise information into non-causal

⁴⁸ Or, less collaboratively, for the dead labour of capital to suck the blood out of the living one, as Marx would put it. For the aims of this section, both perspectives are valid. In Chapter 7, the latter will be preferred.

patterns (Cristianini 2021). This active mediation entangles agencies to the point that, as in all well-performed séances, the sceptic does not know whether they are talking *to* or *through* the medium.

Famously, “a medium is a medium is a medium.” Friedrich Kittler’s sentence stands as a reminder that there is no difference between the logic of occult and technological media (Kittler 1990, 229). A skilled medium, much like the technical apparatus supporting a remote conversation, knows how to disappear behind the *mise-en-scène* and through the staging of other presences (Helen Harper in the dissertation’s prologue performs this role masterfully). And if, as the anecdote of Thales of Miletus shows, movement can be interpreted as the trace of a soul and thus of a disembodied presence, through their apparently autonomous interactions the interfaces I had come up with represented the ideal devices to conjure it. In other words, they allowed me to become a medium myself.

In a pun around the word *Ouija*, the well-known board with letters that two or more people navigate through a planchette to communicate with spirits, and that my colleague Jack Armitage evoked at their first encounter with Stacco (see Section 4.7), I titled the piece *Mouja*, or the *Magnetic Ouija*.⁴⁹ I developed *Mouja* iteratively, and incrementally perfected the narrative of the piece as I developed my skills of “coming and going,” of alternating my own presence, much like in the Voder anecdote, with those conjured by the movements of Stacco’s magnetic marbles.

49 To simplify the title I elided the letter “i” (from *Mouija*, to *Mouja*).



Figure 5.1. *Mouja at Fabryka Sztuki, Łódź. Combining multiple marbles into a pendulum that controls Stacco.*

5.2.2 *Mouja*

*Mouja*⁵⁰ may be divided into two sections. In the first, I perform Stacco on a series of different RAVE models; in the second one, while I let Stacco drone on a particular configuration of magnetic marbles, I perform with Thales, first on the “Icelandic spell” score (see Section 4.7.1), and then with the controllers facing each other. The piece lasts about 25 minutes. To facilitate reproducibility and performance, I composed an embodied sketch with scene transitions and relevant areas in the latent spaces of the models I used (Figure 5.2).

I chose four NAS models for Stacco: a voice model trained at the Intelligent Instruments Lab on a vocal dataset of monophonic audio from singers demonstrating standard and extended

⁵⁰ The video of the performance première is available at <https://zenodo.org/records/20497303>

vocal techniques;⁵¹ the IRCAM Darbuka model, which comes with the default IRCAM nn~ abstraction for MaxMSP, and offers a highly dynamic set of percussive sounds;⁵² a “water” model trained with files from the freesound.org archive featuring liquids flowing;⁵³ a choir model trained by artist Jonathan Chaim Reus (Reus 2023).⁵⁴ Along with the four models, Stacco includes a chaotic FM synthesiser of my design.

The NAS model I used for Thales is the default magnets one (see Section 4.3.3). The piece also features ICST Ambisonics to dynamically spatialise each model in the performative space (Schacher 2010), with each one of Stacco’s attractors controlling the spherical coordinates of one or two models combined into one sound source. Thales’ controllers instead spatialise their respective models, mapping the three-dimensional space to the gyroscope data.⁵⁵

The score (Figure 5.2) develops a clockwise timeline around the attractors, starting and finishing at the top of the board. The timeline is for the performer to follow, with no automatic change of setting. It begins with the voice model, played by moving at a distance a pendulum of magnetic marbles (Figure 5.1). The apparently autonomous movements of the pendulum activate unintelligible speech, setting the mood for the whole piece. As the pendulum moves above the score, the sketch invites the performer to remove the magnetic marbles one by one, redistribute them around the four attractors, and layer the voice with the water model.

After an open improvisation on the first scene, the percussion enter abruptly, with Stacco offsetting timbre transfer values from a drum loop fed into the encoder. Here, the score prescribes adding one more sphere to each attractor, thus making the attractors highly sensitive to each other. Moving a stacked couple of magnets in one thus causes all others to move, increasing the impression of an invisible hand manipulating them.

The piece continues with a cross-fading of the models with two FM chaotic synthesizers, whose high sensitivity to small changes in the magnetic fields (as opposed to RAVE, whose response is relatively slow) shifts the atmosphere: from dynamic, eerie voices and percussive sounds to a minimalist gestural manipulation, where transformations seem to occur with no apparent movement.

51 <https://zenodo.org/records/1193957>

52 <https://github.com/acids-ircam/RAVE>

53 <https://huggingface.co/Intelligent-Instruments-Lab/rave-models>

54 The model was trained on a collection of open choral datasets augmenting core recordings made in collaboration with the Musilon student choir of the University of Twente.

55 Clockwise from leftmost attractor: attractor 1 spatialises FM 1, attractor 2 spatialises Choir and Percussion models, attractor 3 spatialises organ model and FM 2, attractor 4 spatialises voice and water models.

5.2.3 Performance Practice

Among the three works discussed in this chapter, *Mouja* is the one I performed the most. Throughout the different iterations, my performative approach has shifted considerably. During the first exhibitions I typically searched for areas in the latent space that would return realistic sounds or discernible rhythmic patterns, often going back and forth between two or more areas in order to assert a degree of repeatability and control. In hindsight, this approach was too “demonstrative” and rather constrained. But as the practice progressed and my skills and experience with the interfaces and models improved, I transitioned to a more exploratory mode, which enabled a much greater expressivity and creative freedom. This approach may be summarised as a search for surprising and “weird” sounds, and at a higher level, for “weird” combinations of sounds across the different models I layered together.

While Chapter 6 will thoroughly discuss the term *weird* from a hauntographic perspective, here I mean mainly two things by it. First, it refers to what might be described as physically impossible sounds. An effective example of this comes from Philippe Esling, who co-authored RAVE with Antoine Caillon (2021), at his 2024 TENOR Conference keynote. During the speech, Esling overviewed the features of a yet-to-be-released RAVE 2.0, playing with what he described as a “choir with infinite lungs.” To the vocalisations of such a weird choir, *Mouja* juxtaposes the unintelligible chatter in seemingly forgotten languages of the voice model, and layered rhythmic patterns from liquefied percussion.

A second dimension of weirdness, rather than from sounds of impossible bodies with infinite organs, refers to the glitches idiosyncratic to the RAVE architecture, tracing its signature character onto the reconstructed sounds. This is the first encounter with a relevant Derridean notion, which will be further explored in the dissertation in relation to hauntology, that of *trace* as the sign of a deferred presence, which Chapter 6 further develops through the notion of *spectre* (see Section 6.2). In particular, in *Mouja* I became progressively interested in investigating the liminal (here, in the literal sense of being close to the representational boundaries) areas of the latent space where the representation becomes sparse. Here, the effects of digital distortion that are idiosyncratic of the RAVE architecture come to the foreground, turning the reconstructed sounds into otherworldly growls, oceans of digital noise surfaced by the scattered fragments of the training data.

5.2.4 Spectrality in *Mouja*

As practice evolved, so did the response to the piece from the audience. One aspect remained nevertheless constant, that is, the appreciation for the gesturality and bodily engagement that Stacco and Thales require. This connection between gesture and sound became a point of conversation with audiences and curators after many of the performances, and particularly within communities of electroacoustic musicians, who are well aware of the performative and gestural limitations that characterise computer music.

But as my performative skills grew, along with my ability as a medium of disappearing at will behind my interfaces' interactions (a subtle ability of hiding in plain sight, of occupying the stage at some level, but not at *all* levels), audiences began projecting agencies and presences onto them. One telling example comes from the feedback I collected after my 2024 xCoAx performance (Privato 2024). On this occasion, several listeners reported seeing eyes on Stacco's marbles, an effect probably amplified by the stage lights positioned above. This pareidolia effect would probably not have occurred without the seemingly autonomous movement of the magnets, the same leading Thales of Miletus to endow the loadstone with a soul. Along with Thales of Miletus, it is inevitable here to conjure again Jennifer Walshe, and argue that in performing *Mouja*, the question "who is listening, and how?" (2023) came increasingly to the forefront as the core matter of interpretation.

When agency is undefined, it inevitably becomes a vessel for subjective and collective projections. By harnessing this effect, the performance functioned as a probe into the discourse around some of the invisible forces that shape society, at work behind AI's technological enchantment. In the interviews discussed in Chapter 7, this unspecified impression of presence led to conversations about AI as a social agent, and behind it, the role of corporate interests that shape our understanding of this phenomenon. The next piece deals precisely with these issues. Whereas *Mouja* conjures invisible hands seemingly moving marbles on Ouija-like boards, *End of Empire* engages with a very different "invisible hand": that of capitalism itself, as famously described by Adam Smith (1776, 246).

5.3 End of Empire

End of Empire (Figure 5.3) is a large-scale, site-specific interactive installation I developed together with artist Eva Sajovic for the 2023 British Textile Biennial. It deals with themes of

labour extraction and displacement under late capitalism, the relevance of communities in imagining alternative futures, and the role of technology and artistic intervention in times of polycrisis. The installation features four large, hollow textile columns positioned in a wide room, seemingly holding up the roof, with knitted photographs of Greek columns' ruins. Functioning as a large-scale musical instrument, each column independently reacts to changes in proximity and touch by activating and performing a series of real-time NAS models. At the same time, the exhibition space is activated via subwoofers tuned to the resonant frequencies of the hall's structures.

The installation was developed after months of sustained engagement by Eva with the community of Nelson (Lancashire), an area at the centre of the First Industrial Revolution, at the core of which was the textile economy. Through the local In-Situ association,⁵⁶ Eva hosted a series of workshop activities focusing on knowledge resilience in reaction to how consumerism thrives on our ignorance about the origins, materials, and makers of the products we use. In the months that preceded the installation, an intergenerational and intercultural group of residents engaged with Eva in activities of knowledge exchange around knitting, felting, and embroidering, and in sound walks around the area. This turned the process of research and development preceding the realisation of the work into a participatory one.

While this piece would warrant extensive analysis given its scale, technical complexity, and political breadth, here I enter into the specifics of my contribution, contextualising it from a historical and cultural perspective, motivating my artistic and design choices (especially in the use of NAS), and connecting them with the scope of the dissertation.

56 <https://www.in-situ.org.uk/>



Figure 5.3. End of Empire. Eva Sajovic and Nicola Privato (2023).

5.3.1 Context

Nelson, as part of the Lancashire area, was at the centre of the First Industrial Revolution. As such, it represents a unique witness of early capitalism, bearing in its social fabric the traces of a hypertrophic and fast-paced development, attracting capital and workers from all around the world, and in particular labour from the British Commonwealth. Yet as profound as the industrial revolution's impact was in the area, equally profound were the effects of the abrupt crisis that followed the emergence of competing labour markets and the subsequent displacement of the industries.

Walking the streets of Nelson, one is confronted with the spectres of this transition. The canals, excavated to import raw material and export finished goods, would now be completely silent, if it were not for the noise of the highway running parallel, which took their place as the main infrastructure for the transportation of goods; the chimneys of the mills rest without smoke; the metallic structures of the gasometers tower naked through the hills, as ossified

remnants of antique giants; the populous Pakistani community, once for the most part employed in the mills and now active in groceries and open-air markets, seems like a world within a world punctuated by pubs, their regulars visible through the windows as one walks the streets, appearing almost like strangers to their own timeline. If a spectral ethnography can be imagined (Armstrong 2010), Nelson and, by extension, the Lancashire area in which it sits would represent the ideal candidates for such an inquiry.

But where capital brought large-scale exploitation, communities fostered resistance through collective organisation. From workers' actions against the spinning jennies, to the Luddite movement mobilising against machinery that threatened livelihoods, the history of this area shows how capitalist exploitation, class struggle, and technological innovation intertwined throughout the last three centuries. It is through these struggles, and out of early forms of resistance, that communities learnt to organise themselves into the labour movements that would drive much of the twentieth century's social progress within the imperial core.

5.3.2 *AI and Empire*

Looking at scholarly work around the theme, it is increasingly evident how the notion of *empire*, which this installation incorporates in the title, represents a relevant interpretative key to the role of AI in our society. In *Empire of AI* (Hao 2025), Karen Hao draws a parallel between colonial history and the exploitative use of AI technology, identifying a series of key similarities between the two: first, the corporate seizure of social resources, and the re-interpretation of rules to justify it (for instance, the extended appeal to "fair use" clauses); second, the resort to ghost labour from the Global South to clean and label data; third, the control of knowledge-production that derives from the technology's use; and fourth, the production of narratives to justify the cultural hegemony.

From a historical perspective, Pasquinelli offers a materialist account of AI that is grounded in capitalism's automation of labour (Pasquinelli 2023). Rather than adopting the cybernetic framework that tends to dominate scholarly discourse (for instance, see Hui 2023), Pasquinelli draws a parallel between AI and Babbage's Difference Engines. These systems were designed out of precise capitalistic needs, such as automating the repetitive intellectual work necessary to the colonial enterprise. This historical angle ties into the present, for instance, through West, Whittaker and Crawford's work (2019), focusing on how AI systems

reproduce the systemic biases of the corporate groups that design them, a pattern that inevitably also invisibilises subaltern groups (Thompson 2020).

The literature tying imperial interests with those of the corporate groups by which large language models, text-to-image and text-to-audio models are developed is dense and convincing. But if AI serves extractive agendas, how can we begin to imagine a different use, one that serves rather than disempowers communities and fosters connection rather than alienation?

5.3.3 Methodology

My work on *End of Empire* focused on two pillars emerging from these considerations: first, the role of communities in times of polycrisis; second, the possibility of a non-extractivist use of AI, one that contributes to the collective awareness upon which action against exploitative conditions has historically been grounded.

I approached these objectives by organising a series of workshops in the community of Nelson, aimed at collectively collecting audio recordings used to train the AI models.⁵⁷ The groups taking part in the workshops were invited to join sound walks along the canals, the mills, the city centre and the parks, and to collect field recordings from subjectively meaningful places in the area. This involvement was not limited to the workshops, but continued through the subsequent phases, and particularly as I worked on the curation of the datasets I trained the models upon.

Once the material was gathered, it became clear that the recordings required extensive treatment to remove unwanted noise, which could potentially compromise the training. However, the process of cleaning audio by removing hiss, wind noise, cross-talk, cars passing by, while often necessary to obtain a workable dataset, inevitably raises methodological and ethical questions. How does one decide what is relevant and what is disposable? The risk, here, was that of erasing voices, perspectives, and entire narratives through the subjective filter of my own listening.

⁵⁷ Much attention has been dedicated to the ethical issues around the use of the data, and to address power asymmetries between myself, the institution I worked with, and the individual members of the community. For instance, I provided informative material around the use of the data we collected, invited voluntary participation rather than actively searching for volunteers, and insisted on the locality of the work when offered to tour it.

I thus invited participants to listening sessions where we went through the recordings and talked along with the playback, discussed the reasons behind certain choices, and the stories behind those reasons. Through these dialogues, a kaleidoscope of personal narratives emerged: stories of social bonding and struggle, personal and generational memories around migration, integration, and labour, unveiling the connections between Nelson's auditory landscape and the lived experience of the participants. Thanks to these conversations, I could identify which sounds were relevant to the participants and which ones they considered unwanted or irrelevant to the recordings, thus improving the quality of the outcome. But even more importantly, in the process I built a different awareness of the material, which rather than as raw data to mould into a model I could now understand experientially as a carrier of meaning, memory, and identity. This provided the necessary groundwork to inform *End of Empire's* sound and interaction design, ensuring that the work was conceived with, developed through, and enacted within the same community.⁵⁸

This principle of locality also shaped the management of the data and the models. While artists may welcome the circulation of models trained on their playing style across the artistic community, in the case of a participatory and site-specific work such as this one, this circulation would be inherently extractive. For this reason, as agreed with all participants prior to the data collection, after the exhibition the datasets and the models trained on the material they provided were deleted. I view this deletion as a constitutive aspect of this work's ethos.

5.3.4 Sound Design

As a site-specific work, *End of Empire*⁵⁹ confronted the exhibition space both from an architectural standpoint, with the structural trompe l'oeil of the textile columns conveying an impression of support, and from an acoustic one. The latter was achieved by hanging speakers inside each of the four columns in the main room, and by placing four subwoofers in the smaller side rooms, two per side (Figure 5.4).

The subwoofers in the side rooms generated a series of sweeping bass sounds tuned to the resonant frequencies of each of the rooms' large wooden and glass doors (variable from 46

⁵⁸ This non-extractivist, community-focused approach reflects onto the technical means as well. The use of (relatively) small-data models such as RAVE has been shown to limit bias, reduce energy use, and grant a higher transparency both in data collection, attribution, and deployment (Bryan-Kinns et al. 2025).

⁵⁹ A video walk-through of the installation is available at <https://zenodo.org/records/20497396>

to 52 Hz). The sweeping sounds were intermittent and stochastic, sometimes moving circularly through the columns, and sometimes bouncing from one side of the room to the other. When active, the structures of the rooms vibrated sympathetically, conveying a low, deep rumble moving from the walls on the side to the viewer's ears and chest. This artifice functioned as a metaphorical gate granting access into a space where the boundaries between human and architectural bodies dissolve.

Each of the four central speakers embedded in the columns featured a generative model based on the field recordings collected by the local community during the workshop activities. Following the conversations with the participants, I organised the data into four thematic areas, assigning each group to one column (see Figure 5.5).

Column one: The largest, central column featured a RAVE model trained on people's voices as recorded in the sound walks. The model performed both timbre transfer and latent manipulation. For the timbre transfer, I chose a political speech that multiple participants had independently recorded during the town's fair, in which a member of the socialist party addressed the audience with heartfelt tones on workers' rights. Physical interactions of the viewers with the columns were instead used to offset the latents in the decoder's input. The outcome was a quasi-unintelligible speech, with the column eerily addressing the viewers with emphatic tones and a marked British accent, shifting in timbre as curious visitors touched it or stepped inside. The column acquired a solemn, messianic quality, as if bearing an imperative but indecipherable message, fading into an indistinct mass of familiar voices searching for shared words. One, eerily persistent word got nevertheless through. Comrades.

Column two: The second column hosted a RAVE model trained on the singing of birds in the area. Many of the field recordings from the workshop participants consisted indeed of birds' sounds collected all around the town. In Nelson, the canal built to carry cotton and textiles in and out of the mills runs parallel to a highway, whose noise clearly came through in most of the recordings. Far from reinforcing idyllic aspirations of finding refuge in the lost paradise of a world freed from technology, the column's reconstruction of the recordings underlines the fictitiousness of the nature-nurture dualism, and the human interdependence and co-evolution with more-than-human systems. This effect was achieved by offsetting the models' latents into areas where the bird songs acquired a digital nuance, or fused with the traffic noise, and by mixing this body of sounds in-between the organic and the machinic with that of the original dataset.

Columns three and four: I composed the soundscapes of these two columns so that they would engage with each other discursively, each reflecting a different mode of textile production. The third column featured sounds from industrial machinery collected in Nelson's active mills, whereas the fourth one featured sounds from Eva's craft process of dyeing and weaving the textiles used for the columns. These included recordings of the popping of seeds, their boiling and the back-and-forth movements of the loom's shuttle. The two columns, placed three metres from each other, entered into a dialogue foregrounding the sonic differences (and at times similarities) between craftsmanship and mass-production.

Even though the initial idea was to embed one RAVE model per column, I had to come to terms with two limitations: first, the modest amount of processing power of the machine I was provided, which with two models working at the same time was already running close to its CPU limits; second, RAVE's suboptimal results in dealing with datasets featuring a lot of noise (especially from the machinery of the mills), which in newly trained models often converged into homogeneous walls of variably filtered hiss and hum. While I used RAVE in columns one and two, for columns three and four I manipulated the data with simple DSP and granulation techniques, fine-tuning a series of states to be triggered by the viewers' interactions with the columns.




Figure 5.4. End of Empire's technical details. Top-left, column 1 from the inside, with the embedded speaker; Bottom-left, subwoofer activating the space from the side rooms; Right, jumper wires and microcontroller embedded in the textiles, viewed from the inside of column 1.



Figure 5.5. Columns 1, 2, 3, 4.

5.3.5 Interaction

As a work in movement, *End of Empire* invited the viewers to activate and actively engage with the columns (Figure 5.6). Along with a speaker hanging from the ceiling, each column embedded two types of sensors: capacitive and ultrasonic. The sensing surfaces consisted of conductive threads densely woven into the textiles, all connected to a microcontroller stitched inside the column.⁶⁰ Each column featured at least four large sensing areas distributed at head height, each connected to a different input in the microcontroller, so that by touching a particular spot in a column, one could selectively trigger certain behaviours.

The ultrasonic sensors were stitched at chest height inside each of the columns.⁶¹ The columns consisted of two textile surfaces, curved and sewn together from a height of around two metres. The gap between the textiles invited the viewers to step into the columns, and once inside, the person's presence and movement activated the distance measurements of the

⁶⁰ ESP32 TTGO T7 V1.3-mini.

⁶¹ HC-SR04.

ultrasonic sensors. Following the approach of a previous exhibition using similar components,⁶² I intentionally exposed the microcontrollers, sensors and wiring, using them to add to the aesthetic value of the piece. By stepping inside the columns, the viewer would witness the tangle of coloured jumper wires connecting the capacitive sensors to the microcontroller, woven with the textiles into a cybernetic sculpture.

Beneath the columns, a hollow wooden pedestal hosted power supplies and a router for a dedicated LAN, transmitting sensor data wirelessly to the computer responsible for sound synthesis via four MaxMSP patches. The viewers entered the room via a small stair in one corner, immediately confronted with the trompe l'oeil of the columns and a cognitive dissonance between their apparent structural solidity and the lightness of the textiles, subtly breathing through air currents like Calder's Mobiles. A small sign invited viewers to interact.

Upon first contact, the viewer activated stochastic vibrations in the walls and sound from the speaker within that column. The sensors responded to the size of the touched area, applied pressure, and movement of bodies within the columns, offsetting in different ways the RAVE models' autoencoder inputs (or, in columns three and four, selecting different granulation files and parameters).

Since the installation had to function independently throughout the month-long biennial, I designed adaptive algorithms that self-calibrated upon daily activation of the work, and a series of selectable sensitivity settings to compensate for humidity-related fluctuations in the capacitive sensors' readings. This requirement added substantially to the complexity of the work, but proved crucial for its reliability.

After the activation, I visited the exhibition space every day for almost a week to fine-tune details and observe how people engaged with the work. Visitors touched and explored the columns with a mix of wariness (who has not been taught not to touch artworks?) and playfulness; some stepped inside them, examining the materials and sensors. A few boldly touched the jumper wires and the microcontroller (which I thus hot glued in place), trying to understand the interaction design; others tried to engage in sonic conversations with other participants.

62 <https://nicolaprivato.com/plantacapes>



Figure 5.6. Visitors interacting with *End of Empire*.

5.3.6 Spectrality in *End of Empire*

End of Empire was a work on collective history and memory. The ruins of Greek columns printed on the textiles evoked the vestiges of an era long gone, thus paralleling one empire with the other. Walking through the city on the way to the installation site, one encounters similarly eerie remnants throughout: skeletons of disused gasometers around the hills, canals now taken over by the local fauna, chimneys all over, standing as ominous oracles about the destiny of all empires.

When collecting informal impressions among the audience, the familiarity of the sounds came through among those who contributed to the recordings. To the staff of attendants and technicians responsible for activating the work every day throughout the month of the exhibition, the work seemed almost alive and temperamental, responding to changes in the environment such as air currents or humidity in ways that were never fully predictable.

Because of all the preparatory work, I was particularly haunted by the sounds of the four columns, gripped by a sense of familiarity and at the same time of estrangement that I did not come across when I worked on *Mouja*. The voices from the central column were pivotal to producing this effect, as algorithmic remnants of the community I had become familiar with in the curation phase. In a way, the collective memory, pulverised and reassembled through the model, had become a spectre to itself, present at some level and yet at the same time completely estranged: whereas in *Mouja* I became increasingly interested in the impossible sounds and the signature traces of the AI architectures, in *End of Empire* I shifted my attention to the training data, and to the lingering presence of the people, places, and stories infiltrating the model's output.

Indeed, the models of *Mouja* felt like black boxes to be explored. To use the oceanographic metaphor introduced in Section 2.3.5, like oceans to step into and play with the coming and going of the waves or to explore in search of weird creatures. The presences I found in these depths felt as alien as a siphonophore can be. But as I worked on *End of Empire* through sound walks, listening sessions, conversations and sharing, the ocean had become populated by the sunken relics, familiar and yet eerily alien, of the people and places I had encountered.

5.4 Mouja+

Mouja+, the third and final piece I discuss in this chapter, is a performance piece I premiered at SWRL festival, Latvia, in 2024.⁶³ Whereas *Mouja* and *End of Empire* focus on real-time synthesis with small-data architectures, *Mouja+* (Figure 5.7) is based on a performative system that complements real-time NAS with a text-to-audio model in a live setting (see Section 2.2.1), prompting the latter with a series of text scores. Such models (e.g. Udio⁶⁴ and Suno AI⁶⁵) have recently come to the fore due to their capabilities, growing diffusion, and ethical implications. I consider this reflection-through-practice necessary in order to expand the insights I gained with *Mouja* and *End of Empire* beyond small-data NAS.

⁶³ <https://swrl.space>

⁶⁴ <https://udio.com>

⁶⁵ <https://suno.com>

While as discussed in the latter part of the chapter *Mouja+* ties in with the previous works through the lens of spectrality, it also offers three secondary contributions. First, this work pushes the performative boundaries in the use of text scores with NAS models in live settings, a small but emerging area of musical experimentation. Second, it reflects on how the use of language needs to shift as we interact with these models, and develops a series of strategies in this sense (see Section 5.4.5). Third, building on insights from the latter contribution, it offers a reflection on how accessibility and openness of the data positively impact our interaction with generative systems (see Section 5.4.6). This final consideration parallels and complements *End of Empire*'s insights on the artistic and ethical value of a first-person engagement with the data, which scholarly literature increasingly sustains (for instance, see Schroeder and Reuben 2026).

Due to the breadth of the investigation, some theoretical groundwork is necessary in order to situate *Mouja+* in continuity with existing practices, beginning with the intersection of text scores and generative AI.



Figure 5.7. *Mouja+* at SWRL Space. Riga, Latvia (2024).

5.4.1 *Text Scores and AI*

Among contemporary artists experimenting with deep learning and text scores, Jennifer Walshe stands out as one of the most prolific. From 2017 to 2021, Walshe collected a massive dataset of approximately 3,000 text scores, which may be used to train AI models to generate new text scores (Walshe 2021). This growing dataset contains scores from Oliveros, George Brecht, Mieko Shiomi, and others, but is also open to the contribution of any artist willing to send their work.

Whereas Walshe's project focuses on the algorithmic generation of text scores, leaving to the human performers the role of enacting them, the recent emergence of NAS techniques has introduced new possibilities. Tools such as Udio, Suno AI and Stable Audio (Evans et al. 2024) are now capable of synthesising coherent sounds and musical structures from textual descriptors, allowing composers to design text scores to be interpreted by generative systems.

An example of a performative use of text-to-audio NAS is that of Dadabots' prompt-jockeying live sets.⁶⁶ In prompt-jockeying, music is generated onstage by the NAS model based on live-prompted textual cues. According to Dadabots, "if a DJ must be fluent in their Rekordbox collection, a PJ must be fluent in their PyTorch model."⁶⁷ Such fluency presupposes high programming skills, but also many hours of experience in prompting a particular model. Working with textual instructions on a generative system is indeed very different from DJing, and, what interests us most, from composing text scores for human performers. As we shift to the NAS model, our use of natural language needs indeed to change and adapt.

We can look at this from at least two perspectives. First, we can look at it from the angle of the implicit, culturally dependent and contextual rules upon which communication is based, and how these rules shift as we address a non-human agent. This aligns with Wittgenstein's notion of language games, according to which the meaning of a word is never fixed but emerges through contextual use (Wittgenstein 1968). Second, we can look at how the potential for expression of text scores changes when they no longer instruct actions to real bodies onstage, but are rather used by an AI model to reconstruct the sonic imprints of bodies and actions that never were.

⁶⁶ <https://dadabots.com/>

⁶⁷ <https://iil.is/news/dadabots-workshop>

The enactment of Yoko Ono’s *Stone Piece*, from the *Fluxus Performance Workbook* (Friedman and Smith 2002), offers a fitting example of the former case. The piece invites the performer to “take the sound of a stone ageing.” This is an ambiguous statement on many levels. What is the sound of ageing? How does a stone age? And as a consequence, what does ageing mean in this context? Like a haiku, this score introduces interpretative ambiguities that call for introspection. Yet, to a question requiring a certain awareness of the language game being played and perhaps to consider less simplistic interpretations, the model might (and indeed did, as I tried the text with the pipeline described below) default to the sound of a stone being scratched.

Whereas Yoko Ono’s piece introduces a problem of interpretation, Jed Curtis’ provocative *Music for Wise Men* prompts reflections on disembodiment and signification. The score instructs performers to “commit suicide” (Friedman and Smith 2002). These instructions inevitably confront a human performer, with a physical body they are bound to, with the limits of literal execution of the score. However, when an AI interprets the same instruction, the absence of a body prevents confronting the audience with the foundational paradox of Curtis’ score; at the same time, from this (at least for Curtis) unforeseen rearrangement of material affordances, new potentials for interpretation can emerge.

These initial considerations prompted this investigation into the possibilities and limitations of composing and performing with text scores and text-to-audio NAS. In *Mouja+*, which I introduce in the next section, I focus purposely on scores from the Fluxus collective. By experimenting with notation meant for human actors, in which sound is seen as a consequence of a certain performativity, I could bring to the forefront the model’s interpretative limitations and use the findings to inform compositional strategies and reflections concerning text-to-audio NAS.

5.4.2 System Design

Mouja+ leverages Stable Audio Open (Evans et al. 2024), a text-to-audio model trained on the Freesound archive.⁶⁸ On the client side, a Graphical User Interface (GUI) communicates with a remote server through an OSC application exposing Stable Audio API endpoints.⁶⁹ By

⁶⁸ <https://freesound.org/>

⁶⁹ <https://github.com/elgiano/gradio-osc>

processing the prompt on a remote server with a high-performing GPU, I could significantly reduce the time required for the generation of the sound files (around 20 seconds for a 47-second sample) compared with most consumer machines, which is a critical aspect during live performance.

On the client GUI (Figure 5.8), one writes positive and negative text prompts,⁷⁰ and defines length in seconds, number of steps, temperature,⁷¹ and buffer for the download. The client forwards the prompt to the server, which generates an audio file. Once ready, the file automatically uploads into the specified buffer on the client computer. The GUI allows one to view, play, process and mix up to eight audio files at the same time, which can be routed either to the speakers or to the real-time NAS model (RAVE) for timbre transfer.

When timbre transfer is applied, control signals from Stacco are added to the RAVE latents before the decoder, allowing one to influence the resynthesis with Stacco. The RAVE models used for timbre transfer are all trained at the Intelligent Instruments Lab. All models have been collected from open-source material or with the explicit consent of all parties involved, are free to download and use, and comprise foley sounds, liquids, guitars, voices, percussion, magnets, organs, and birds singing.⁷²

The pipeline also applies Ambisonics spatialisation, since *Mouja+* was premiered on a 36+2-speaker dome (Figure 5.10). The readings from each of the four magnetic attractors in Stacco are mapped to the spherical coordinates of one sound source within the Ambisonics space using ICST Ambisonics (Schacher 2010). The movement of the four sound sources is automated, but at all times it is possible to switch to manual control with Stacco.

70 Negative prompts are explicit instructions on what content an AI model should avoid in its outputs.

71 In generative models, temperature defines the amount of randomness in the model's output.

72 The models are available at <https://huggingface.co/Intelligent-Instruments-Lab/rave-models>.



Figure 5.8. Mouja+ GUI. The top section handles text prompts, the left one playback of the generated files, the right one timbre transfer and latent manipulation with RAVE.

5.4.3 Structure

Mouja+ builds on *Mouja* by complementing it with a series of text scores adapted from the *Fluxus Performance Workbook* (Friedman and Smith 2002), prompted onstage to the text-to-audio system, and whose output I manipulate in real time with my interfaces. In this expanded version, the circular timeline around the embodied sketch that describes the different scenes in *Mouja* (Figure 5.2) is compressed into a shorter solo with *Stacco*, which I play in the middle of the piece.

The Fluxus scores, ten in total, are displayed for the audience on a large screen, along with information on whether the synthesis from each successive prompt is in process or complete (Figure 5.9). The structure is the result of a careful selection and readaptation of the ten scores, overlapped and juxtaposed into a tripartite form. Between the first and second section I play the *Mouja* solo on *Stacco*, and between the second and the third, I play the *Mouja* solo on *Thales*.

The first section consists of layers of concrete sounds moving concentrically around the audience, and begins with the sound of a heavy resonating body (a piano) being dragged on the

stage, and whose position in the Ambisonics space is controlled using Stacco. The text scores are three: “Let piano movers carry the piano on stage,” “Wash the piano, wax it and polish it well” (adapted from Nam June Paik), and “Eat a juicy apple” (adapted from Bob Lens). As the files are generated, I use Stacco to control transitions and dynamics. The first section ends with a solo on Stacco, using direct latent navigation and no resynthesis, based on the score in Figure 5.2.

In the second section, the concrete sounds become pointillistic. The scores are “Hammer nails into the piano keys” (adapted from Tomas Schmit), “Dropping coins” (excerpt from George Maciunas’ *Solo for Rich Man*) and “The sound of the stone ageing” (adapted from Yoko Ono). As in the first section, I resynthesise the sounds using Stacco’s embodied sketch. This second section ends with a solo on Thales, as in *Mouja*’s latter part. The third section conjures onstage an open market, with unintelligible voices emerging from a crowd, which I dynamically displace in the Ambisonics space. Finally, the sounds deflate again into that of a piano being dragged off the stage. The text scores I use in this section are “Performer sells the theatre” (adapted from Ben Vautier), “Piano movers carry piano out of stage” (adapted from Nam June Paik), and, as a conclusion, Ono’s “The sound of the stone ageing.”




Figure 5.9. *Mouja+* score display.

5.4.4 Reception and Evaluation

I performed *Mouja+* at the SWRL 2024 festival, in Riga, Latvia (Figure 5.10).⁷³ After the performance, I briefly explained to the audience Stacco's workings and *Mouja+*'s technical pipeline. While time constraints prevented extended discussion, I gathered informal feedback during the following days of the festival. Among the insights I collected, two are particularly relevant here, and will be used to ground the discussion in the next sections.

First, the visual emphasis on the model's real-time generation of the sound files brought to the forefront the agency of the text-to-audio model, effectively incorporating it as part of the performance. The second insight concerns instead the model's interpretative failures, which caused some cognitive dissonance in the audience between the expectations from the textual instructions and the sonic outcomes. One audience member, for instance, commented on how the model reconstructed the hammering of nails onto the piano keys in Schmit's piece with no consideration of the word "hammering," resulting in dense, minuscule sounds of nails falling on a surface instead of the sharp and metallic thump one would expect from a hammer, accompanied by sparse and feeble notes from a piano. This clearly exemplifies the tension identified in Section 5.4.1 between human and AI's understanding of language.

In hindsight, from my perspective as a composer and performer of this piece, two reflections are worth offering. First, while the technical pipeline described in Section 5.4.2 worked as expected, the setup, combining experimental NIMEs with remote text-to-audio prompting and spatialisation in an Ambisonics dome, was elaborate to the point that it could not be reproduced in most live situations. This was a relevant limitation, since I could not develop a sustained practice as I did with *Mouja*.

To expand the comparison between the two pieces with a reflection on my own role, it is also worth noting that in *Mouja+* I performed more as a laptop musician than as an instrumentalist, whereas in *Mouja* I performed primarily as an instrumentalist. Coming from a jazz background, I feel more at ease when performing an instrument than when prompting text at a computer. While this piece was crucial for the research in that it allowed me to expand the investigation beyond NAS, if I were to use this pipeline again, it would be as a minor complement to a performance with Stacco or Thales, rather than as the core structural element.

73 Audio recording of the performance may be accessed at <https://zenodo.org/records/20493761>



Figure 5.10. Performing Mouja+. Solo on Thales before the third section.

5.4.5 Compositional Strategies

As I worked with the Fluxus scores, I had an initial idea of what sounds I was looking for, how much variability in the generation of the files I was willing to accept, and how I intended to combine them to fit a metanarrative. At the same time, I had to be aware of the interpretative limitations of the model, and thus developed a series of strategies for mitigating them. Of these, three proved particularly useful.

Dataset exploration: Stable Audio Open is trained on the Freesound archive, which is open and freely accessible online.⁷⁴ I used the Freesound advanced search engine to explore text embedding, typing keywords and selecting tags based on the scores I was experimenting with. One example is the word “exit.” In Freesound, this word is often associated with environmental recordings involving trains arriving at stations, opening and closing their doors, and leaving.

⁷⁴ <https://freesound.org/>

This suggests that the statistical likelihood that prompts containing this word might incorporate sounds of moving trains is high.

Strategic use of negative prompts: Once I understood how words were embedded in the original data, I typically decided whether to discard a score or use negative prompting to steer the system towards a desired outcome. In most cases, negative prompting was the most effective strategy. For instance, useful negative prompts when using the word “exit” were “train” and “bus,” “station,” or “soundscape,” which allowed me to avoid the sounds of public transportation coming and going.

Prompt micro-tuning: As I tested different text scores, I realised that by slightly changing the grammar and by inverting the order of the words, I could guide the generation of the sounds towards a desired outcome or make a certain output more consistent through successive prompts. “Let piano movers carry the piano out of the stage,” “piano movers carry piano out of stage,” or “eat juicy apple(s) during concert” and “eat a juicy apple during concert” tend to be interpreted in similar ways by humans. However, by introducing small changes, such as by adding “a” and removing “(s)” from the text score, I could fine-tune the generation, or increase the likelihood of generating a certain result given a certain prompt.

These three strategies were critical in adapting my use of natural language to Stable Audio’s interpretative constraints. Nonetheless, the possibility of the model generating unexpected and sometimes undesired outputs onstage was always present. For instance, as noted in Section 5.4.4, the hammering of nails on the piano keys in the text score adapted from Schmit’s piece was hardly the sharp hit produced by a hammer hitting a nail, which inevitably defused the intention of the original text score.

5.4.6 Dataset Exploration and Open Access

Of the three methods, dataset exploration deserves further consideration. Understanding the relationship between individual keywords (or clusters of words) and the dataset through the search engine offered by Freesound turned out to be an effective strategy in designing text scores for Stable Audio Open, one that I iteratively applied to every score I worked on. Crucially, the applicability of this method relies on the openness of the Freesound archive and

on the availability of an advanced search engine, two features that are not available in most text-to-audio models.

Indeed, the academic literature pointing to the value of open access to datasets is rich. For instance, FAIR practices (Findable, Accessible, Interoperable, Reusable) to data collection and deployment are increasingly adopted in order to mitigate bias (Raza et al. 2024). Moreover, the ethical relevance of data transparency is being increasingly discussed (Hardinges et al. 2023), along with the systemic lack of data documentation among AI companies (Longpre et al. 2024).

In response to these pressing issues, researchers have implemented tools to trace lineages and audit text datasets (Lewis et al. 2020) and multimodal content, for instance allowing one to query a certain corpus, to assess whether and how data is licensed, or assess diversity, inclusion, and biases. Here, I tackled this issue from a slightly different angle, showing how, instead of relying on tentative prompting approaches, we can more effectively refine our natural language instructions through a direct engagement with the training data, provided access and effective tools to query it.

Open access and accessibility are thus essential not solely in addressing critical issues around bias, data ownership and accountability, but also to facilitate our engagement with text-based generative systems.

5.4.7 *Spectrality in Mouja+*

As an expansion of *Mouja*, *Mouja+* reflects on spectrality and AI as well, but addresses the theme from quite a different perspective. Through the use of scores that incorporate concrete objects and human actions (such as piano movers, or the final selling of a theatre which conjures an entire crowd on the stage), their projection for the audience through a descriptive score, and the active displacement of the sounds on the stage by means of spatialisation techniques, *Mouja+* engaged the audience in a playful conjuring of imaginary bodies and actions.

The sounds of dragging pianos, the chewing of fruits, the clinking of coins, or the murmuring from the crowd bore witness to things in the world they seemed to represent, but never were because generated in the moment. This space of ontological uncertainty opened up new interpretative dimensions in the text scores for me to play with. The act of chewing an apple, for instance, became subtly uncanny, seemingly performed by a ghost walking among

the audience (we will see in Chapter 6 how, according to Derrida, ghosts exist as *work* without presence). Similarly, the piano dragged on the stage became as present (or absent) as the piano movers that dragged it, and the conjured crowds almost tangible while bearing in their unintelligibility the traces of their algorithmic origin. The more it was made evident (e.g. through the use of a descriptive score and spatialisation techniques) that the generated sounds did not signify things in the world albeit appearing as testimonies of real bodies and actions, the more their paradoxical referentiality was amplified along with the twisted causality it bore.

This ontological uncertainty opened up interesting expressive possibilities, but also rendered certain others inaccessible. Especially those in which meaning relies on physical presence. Maciunas' *Piano Piece n.13* prescribes nailing the keys of a piano until it produces no sound, a politically charged act given this instrument's symbolic tie with Western musical history. Importantly, the power of this act lies precisely in its materiality and irreversibility. What remains of this radicality when only the idea of a piano is brought on stage and the critique can be iterated ad infinitum without hammering a single nail?

5.5 Haunted Work

The chapter's artistic reflections presented different variations around a theme, which also emerged in the design phases of Thales and Stacco. Namely, that of spectrality as defined in Section 2.5. *Mouja* focused on the traces impressed by the generative model onto the reconstructed sounds, and on the audience's propensity to project onto Stacco's magnetic action presences of different kinds. *End of Empire* focused on the traces of the local community impressed in the generative model, and on how I became aware of them by engaging in data collection and conversations within that very community. *Mouja+* focused on the relationship between the AI-generated sounds and the "ghostly referents" they conjured onstage.

Before drawing on these insights to address the first research question that this project asks, concerning the discourse and metaphors that emerge as we engage artistically with AI, it remains to be established whether the themes discussed above emerge or do not emerge consistently in the works of other artists. In the next section, I overview a series of works to make the case for the former hypothesis.

5.5.1 *Chunks, Feathers, Mirrors*

While many of the interfaces developed at the Intelligent Instruments Lab are endowed with a certain agency (I think for instance of Notochord’s conversational disposition (Shepardson et al. 2022) or of the Halldorophone’s capricious feedback loops (Úlfarsson 2018)), one interface stands out as particularly kin to the ones I developed in the way it plays with ontological uncertainty. Namely, Giacomo Lepri’s *Pluma*, a sonic sculpture using birds’ feathers painted with conductive ink, which activate and control an FM synthesizer in response to proximity and bodily interaction. The piece reacts to touch and physical proximity by fluttering and flickering, an animal mimicry that is achieved through a mechanical transducer mounted on the body. As Lepri argues, *Pluma* plays with the idea of AI without AI, thus exploring ways to evoke an illusion of liveness and presence (of AI, of an animal, in any case, of some-one or some-thing) in the act of interacting with the digital artwork.

Algorithmic presence (or the illusion thereof) is also investigated by Palle Dahlstedt’s work on *minimal algorithms* (Dahlstedt 2025), relatively simple systems designed to query how small an interactive system can be to evoke the feeling of “somebody there.” The ontological ambiguity that Dahlstedt explores comes from the uncanny recognition of one’s cognitive processes reflected in computational form, externalised and estranged enough to become alien.

Jennifer Walshe and Memo Akten’s *Ultrachunk* (2018) foregrounds a similar ambiguity, though with an additional temporal and mnemonic dimension. In this performance piece, Walshe duets with a model trained on chunks of her voice recorded in successive daily sessions throughout a year. As Akten puts it, in this piece “the original and virtual Walshe inhabit the uncanny valley together.”⁷⁵ Where Dahlstedt’s *minimal algorithms* leave us uncertain as to whether the model’s traces are imprinting on the performer or vice versa, *Ultrachunk* immediately confronts the viewer with the performer’s past self as a spectral interlocutor. Yet, as the piece progresses, the algorithm’s interpretative gaps, glitches, and inconsistencies take over. These “failures” amplify the sense of alienness, transforming Walshe’s past self into a sort of doppelgänger of algorithmic mediation. Fisher offers an ideal synthesis of both Walshe’s and Dahlstedt’s approaches to spectral presence, when he aptly writes that “there are ghosts in the machine. We are they, and they are we” (Fisher 2016, 109).

⁷⁵ <https://www.memo.tv/works/ultrachunk/>

Of a more eerie nature is instead Holly Herndon's *Proto*, an album featuring *Spawn*, an AI model trained by the artist through private and public singing sessions.⁷⁶ In the video of *Godmother*, one of the album's pieces, Herndon merges oblique vocal reconstructions, digitally distorted percussive sounds, and close-ups of the artist's face as viewed by a machinic observer.⁷⁷ While the piece nods at Shelleyan narratives of artificial life, the sense of presence it foregrounds emerges from another ambiguity as well: as we listen to the artist's AI-mediated voice, and as we follow the video which places us inside a machine as it experiences its human interlocutors, we are left wondering whether we have become machines ourselves.

The examples discussed above, namely Lepri's *Pluma*, Dahlstedt's minimal algorithms, Walshe's *Ultrachunk*, and Herndon's *Proto*, are examples of a spectral tendency that extends beyond my individual practice. While the list of such works could be extended considerably, the pieces presented here should suffice to establish a convergence of AI-mediated cultural production toward exploring similar questions and through similar metaphors.

5.5.2 *The Open Work Gets Haunted*

This chapter's investigation began with Eco's notion of works in movement, whose open and unfinished form allows for the situated and contextual emergence of poetic meaning. The open work deals with the deliberate openness that characterises modern cultural production, a perspective that, as discussed in Section 2.4.1, becomes even more relevant with generative technologies. These introduce a type of indeterminacy and openness through the unfathomed nature of their hidden layers, whose depths artists enjoy exploring and inhabiting as a space of open potential.

Here is where Eco's framework begins to intersect with Derrida's critique of the metaphysics of presence and attention to the deferring nature of meaning. Derrida's discourse on hauntology, though, expands this deferral into the territories of presence, which in itself becomes an indeterminate and ambiguous dimension, an open potential that exists only in a movement of exorcism (see Section 6.2). As Chapters 6 and 7 will further show, this

⁷⁶ [https://en.wikipedia.org/wiki/Proto_\(Holly_Herndon_album\)](https://en.wikipedia.org/wiki/Proto_(Holly_Herndon_album))

⁷⁷ <https://www.youtube.com/watch?v=sc9OjL6Mjqo>

spectralisation of presence becomes particularly apparent with media technologies, and even more so in the case of AI.

Based on these considerations, on the insights I gathered through design and artistic practice, and on their expansion into other artists' work, I am now able to respond to the first question that this dissertation asks, concerning the discursive metaphors emerging with our use of AI, and in particular as far as AI's opaque nature is concerned. I argue that the hauntological dynamics that Derrida ascribes to media technologies become particularly pronounced in AI-mediated artistic practices, signalling the emergence of a new metaphor within the epistemological space opened by Eco, one that may be tentatively described with the phrase *haunted work*.

The notion of haunted work does not reject Eco's far-sighted observations on openness and indeterminacy in modern cultural production. It rather focuses on how, with the introduction of generative technologies, entropy and emergence tend to coalesce around questions regarding deferred presence. In the haunted work, the liminal becomes a generative force, even a performing subject, and poetic meaning arises from the traces of what is removed, excluded, deferred, hidden or repressed, and yet haunts from a place of open potential (e.g. the traces of the model-as-agent, of the sources, of referents that never were, as respectively outlined in *Mouja*, *End of Empire*, *Mouja+*). While works in movement purposely include entropy and indeterminacy, haunted works such as *Mouja*, *End of Empire*, and *Mouja+* build on an additional dimension of indeterminacy, a conjuration of deferred presences haunting from within, around, before and after the machine.

5.5.3 Co-creation, Performativity, Haunting

The expansion of openness and indeterminacy into the territories of presence that the haunted work describes inevitably raises the question of what position the composer comes to assume. While Chapters 6 and 7 further develop notions of agency and co-creation from a hauntological angle, it is nonetheless worth considering these questions locally, through the works presented in this chapter. To do this, I begin by distinguishing between the performative use of metaphors dealing with presence, co-creation and haunting, and functional accounts of co-creation with the model-agent outside the performance.

Let us begin with *Mouja*, by adopting a functional account. As anticipated in Section 4.2.2, real-time generative models such as RAVE locate co-creation primarily in the

compositional phase, which involves the curation of the data and the training of the model. The composer outsources part of the process (the training) to the system, and once the model is trained needs to adopt an exploratory approach (see Section 5.2.3). This is why the embodied sketching method described in Section 4.4.3 proved so effective for composers working with RAVE. From a performative perspective, though, *Mouja* does not conjure onstage the kind of agent described above. Rather, it summons a disembodied presence that operates as a co-creative partner in the fiction of the work: here, the aim is to foreground the same contradiction of automatic autonomy described by Derrida (1994a, 192) and introduced in the dissertation's prologue (see Section 1.1), and from there open a space for reflection on AI as a social and cultural agent.

In *End of Empire*, the performative and functional accounts partly overlap. Even small-data models such as RAVE typically require a non-negligible amount of training data. Thus, the phase of data collection becomes a crucial site of co-creation, potentially involving other people besides the composer. This process can be extractive, collaborative, or anything in between. In the case of an extractive approach, the sources involved in the creative process are invisibilised or erased; a collaborative approach instead views the model as a means to organise and access archival memory in unforeseen ways (for a hauntographically informed discussion on this, see Section 7.3), and thus bears witness to the agency of the sources, who become co-creative partners. In *End of Empire*, the community that contributed to the recordings co-authors the work and, at the same time, constitutes the very presence that haunts its outputs (see Section 5.3). This convergence, although not perfect in that the model's mediation is also responsible for the spectralisation of the community's voices, offers a strong argument for adopting data collection and AI model training as participatory practice, provided that a clear ethical stance is assumed.

In *Mouja+*, the performative and functional accounts diverge again. Text-to-audio generative models such as the one used in *Mouja+* introduce an element of co-creation *within* the performance. The statistical variability in the generation of outputs introduces a certain unpredictability, which also needs to be accounted for in the composition phase where one develops and sharpens textual instructions. As Section 5.4 has shown, further variability is introduced by the inherent limitations of the model in interpreting context. Yet in *Mouja+* the co-creative agency of the model is performatively conjured only insofar as it bears witness to the impossibility of the presences signified by the scores, the “real” ghosts at the centre of this haunted work.

The three works described in this chapter feature localised areas of co-creation that do not necessarily erode the composer's primacy as the designer of a composed system. Yet as these pieces exemplify, in the haunted work co-creation is also *staged*. As in the case of *Mouja* or in Walshe's question of who or what is listening, this *staging* (a key hauntological and hauntographic construct as we will see in Chapter 6) hides the composer's hand precisely in order to make the question of agency visible. How this haunted performativity extends into the discourse concerning AI is the question that the hauntographic framework developed in the next chapter addresses.

5.6 Haunting in Digital Culture

Before developing the hauntographic framework and as a way of grounding its relevance, it is worth outlining how the question of who, or what, is listening with AI, transcends the performative stage and yet, it might be argued, not performativity tout court. Take for instance the “discovery” of Loab, a fictional character that in a viral Tweet Steph Maj Swanson claims to have encountered in the latent space of a text-to-image model.⁷⁸ As the story goes, the artist accidentally summoned Loab in an unspecified model through the negative prompt “Brando::-1”, that is to say, “show me the opposite of Marlon Brando's face.” Out of the first prompt, the model generated the letters “DIGITA PNTICS”, whose successive negative prompt revealed the image of a “devastated-looking woman with defined triangles of rosacea on her cheeks.”

Loab (interestingly enough, described as a “she” in the digital meme culture) soon was to be found everywhere, seemingly haunting every image she touched, in that every visual or textual prompt attempting to generate another image from one that contained her semblance would inevitably converge back to it. We do not know whether this might be indicative of a self-contained area in the latent space, of biased prompting, or even not true at all since the architecture remains undisclosed. But this indeterminacy is precisely the point. The complexity of the ritual and the supposedly extraordinary behaviour of the model under the given circumstance perform the conjuration, providing Loab with enough mystery to become a ghost, and the ghost with enough agency to surf the digital imagery.⁷⁹

⁷⁸ <https://en.wikipedia.org/wiki/Loab>

⁷⁹ <https://www.wired.com/story/generative-ai-marquis-de-sade-chatgpt/>

Loab is not an isolated case. Another AI cryptid goes by the name Crungus. Similarly to Loab, this imaginary and grotesque creature was allegedly found by a Twitch streamer while prompting its name on the Craiyon⁸⁰ text-to-image model.⁸¹ These and other creatures haunting the latent space of generative models bear witness to something crucial for the purposes of this dissertation: namely, a broader impulse to project presences and intentionality onto processes whose workings remain concealed, to narrativise algorithmic opacity through spectral discourse, and to use such discourse to make sense of (or conjure away anxieties on) these technologies.

The pervasiveness of spectral metaphors in relation to AI necessitates theoretical tools capable of interrogating these phenomena and their reasons. Chapter 6 responds to this need by developing a framework to artistically and critically engage with the spectral dimensions of generative technologies.

5.7 Conclusions

Chapter 4 discussed how insights from user studies, practice-based investigation and instrument design contributed to a shift in this research, which, from an embodied and pragmatic approach to XAI, began to focus on the meaning-making potential of AI's unknowns. While Chapter 4 looked at the plastic and figurative features of the interfaces I designed around NAS, Chapter 5 expanded the investigation into communicative and meta-communicative levels (see Section 3.4.1).

This chapter offered three artistic contributions, each reflecting on spectrality from a different angle. The first work, titled *Mouja*, applies séance aesthetics as a compositional method. As I performed *Mouja* from one venue to the other, I began to increasingly navigate the foldings of the latent space searching for “impossible sounds,” such as choirs with infinite lungs or unintelligible voices, and looking for the traces of the model’s algorithmic ruminations through artefacts of digital distortion, glitches and latent gaps. My improving skills in playing the role of the medium and performatively disappearing behind the *mise-en-scène* were matched by a tendency among the audiences to project presences of different kinds behind the magnetic marbles.

⁸⁰ <https://www.craiyon.com/en>

⁸¹ <https://en.wikipedia.org/wiki/Crungus>

End of Empire shifts the focus from the invisible hands moving my magnets to the invisible hand of capitalism. Here I focused on a non-extractivist use of AI, centred around active participation and community kinship. As the raw material collected during participatory activities was processed and activated by the NAS models, collective memory became estranged while at the same time maintaining an eerie sense of familiarity. This insight will be further developed in Section 7.4 through the notion of *spectrophonia*.

Mouja+ investigates text scores in combination with text-to-audio NAS. The piece introduces a pipeline for real-time NAS complementing audio generated with a text-to-audio model with prompts based on Fluxus scores and embodied sketching via my magnetic interfaces. Here, I reflected on how AI's statistical processing of language introduces interpretative gaps between the scores and the model's output, and proposed a series of prompting strategies to productively constrain (not remove) the indeterminacy introduced by this dynamic.

In *Mouja+*, spectrality manifested *as* textuality. The AI-generated sounds conjured possibilities rather than pre-existing referents, bearing witness to their subjects' radical absence, their *not being there* but also the ontological paradox of *having never been* in the first place. Whereas *End of Empire* haunted from the loci of memory, *Mouja+* introduced an element of haunting from the future, with the sounds conjuring the ghost of bodies and actions that never were.

Building on the different yet intertwined types of spectrality emerging in these works, I drew on Derrida's notion of hauntology to theorise the emergence of a new metaphor within the epistemological space opened by the encounter of Eco's framework with generative AI, which I termed *haunted work*. This notion focuses on how AI-based art concretises algorithmic processes around a deferring, elusive sense of presence upon which meaning is constructed. In the haunted work, the ontologically liminal becomes a generative force, with presence arising from a performative conjuration. This notion extends far beyond my practice, to encompass a growing corpus of AI-based works in which artists deliberately engage with questions of presence and absence as a source of poetic meaning.

Chapters 4 and 5 have presented the insights from practice underpinning Chapter 6's theoretical framework, which aims to answer the second research question: how do the discursive practices I identified contribute to developing tools for interrogating shifting notions of agency, authorship, and intentionality in human-AI co-creation? To develop a vocabulary for interpreting the findings discussed so far and extend the insights into AI broadly understood,

Chapter 6 offers a comparative reading of the two main hauntology theorists, namely Jacques Derrida and Mark Fisher.

6 Hauntography

*Thou art a scholar; speak to it, Horatio.
William Shakespeare (1604).*

6.1 Introduction

“Time is out of joint.” These words, uttered by Hamlet in witnessing the apparition of his father’s ghost in Shakespeare’s play, give the title to one of Philip K. Dick’s novels (Dick 1984). The plot centres around Gumm, a war veteran living in an idyllic 1950s town, who progressively uncovers the secrets behind the artificially constructed reality he inhabits, built on the ruins of his real hometown. Both Hamlet and Gumm are faced with a series of baffling events that overlap the ordinary with the otherworldly. Whereas in Shakespeare’s play a ghost finds its way back to the living to disclose an unspeakable secret, in Dick’s novel it is Marilyn Monroe’s iconic image, unknown in Gumm’s reality, that appears in a magazine to reveal a world within a world. These inexplicable events produce distortions in the characters’ perception of the temporal flow, pointing to something unfinished in the present state of affairs and to a past that refuses to rest.

The phrase “time is out of joint” provides the ideal connecting thread between the two thinkers that this chapter engages with. On the one hand, Jacques Derrida, who in *Specters of Marx* conjures Hamlet’s ghost from the rise of the curtain (Derrida 1994a, 2); on the other hand, Mark Fisher, who in *The Weird and the Eerie* uses Philip K. Dick’s novel to trace two fictional and affective modes of the *strange* (Fisher 2016). As will become apparent in reading this chapter, this shared foundation is indicative of the thinkers’ common engagement with what lies beyond ontology. With hauntology.

The following dive into hauntology responds to the insights emerging from practice. With the notion of the *haunted work*, findings from Chapter 4 on Un-XAI and Chapter 5 on presence in AI-mediated art find a common denominator in spectrality theories’ focus on epistemological and ontological liminality (see Section 2.5), to which hauntology is foundational. However, while the haunted work establishes a common foundation for an increasing number of AI-based works, it does not provide the tools to interpret and analyse their specificities, differences and contradictions, and from there expand the analysis into AI

as a cultural phenomenon as per the third research question this dissertation poses (see Section 1.3). In preparation for the analysis of Chapter 7, this chapter develops a framework to encounter, observe, and analyse phenomena from the perspective of spectrality.

Developing this framework has arguably been the most challenging task I undertook during these three years, for at least three reasons: first, the limitations of my theoretical preparation, which I had to extend considerably as the research unfolded; second, the rhizomatic proliferation of spectral theories across disciplines, including cultural theory, philosophy, media theory, anthropology, and many other scholarly domains, each characterised by shifts in vocabulary, methods, and conceptual emphasis.

The third challenge I encountered corresponds to a further motivation for developing this framework: namely, the ambiguity of the discourse surrounding spectrality theories, particularly as these are put into practice. As the investigation progressed and I engaged in numerous conversations with peers investigating or applying such frameworks, it became increasingly evident that the lack of a shared vocabulary (I think for instance of terms such as *spectres*, *ghosts*, *spirits*, often used interchangeably) and the tendency of scholarly and public discourse to project disparate and sometimes conflicting ideas onto this domain can constitute a significant epistemological limitation.

While the value of spectrality theories derives in large part precisely from the ambiguous ontological terrain they occupy, I contend that developing a stable epistemology of the spectral is nonetheless both possible and beneficial, especially as we search for theoretical tools to iterate from practice to theory, and vice versa. In light of these considerations, which echo Fisher's own frustration with Derrida's tendency to recede into indeterminacy (see Section 1.2), this chapter attempts to offer a vocabulary and framework to discuss phenomena, and AI in particular, from the perspective of spectrality.

The journey toward *hauntography*, the framework I introduce in this chapter, was inevitably punctuated by pitfalls, gaps, and dead ends, which by no means do I claim to have entirely averted. Nevertheless, I do claim that by the end of this chapter, the reader will have gained at least some insights into how spectral phenomena may be experienced affectively and built upon through creative practice. To reach this goal, in the following pages I propose a comparative reading of *Specters of Marx* (Derrida 1994a) and Fisher's broader corpus of works, with particular attention to his last and posthumous ones, and informed by the insights from practice.

I begin by contextualising the two authors historically and showing how both frame their hauntological reflections around a critique of the rhetoric of the *end of history*, an idea

that, while not new, gained momentum at the beginning of the 1990s, after the collapse of the Soviet Union and as the Western bloc transitioned into the neoliberal era. This historical perspective translates for Derrida into a metaphysical dimension wherein all ontologies lie at the end of history, inherently haunted by the seeds of their undoing, and for Fisher into an epistemological tension toward the outside and the new, to which the framework I propose adheres. From this common plateau, I continue examining Derrida's hauntological framing of ghosts and spectres, and observing how his deconstructive stance translates in Fisher's work into a creative one. I then discuss Fisher's affective and fictional categories of the *strange*, the *weird* and the *eerie* (Fisher 2016), from the perspective of hauntological thinking, suggesting that these categories parallel and extend Derrida's reflections on the spectre's *out of joint* and liminal nature, and that these may be viewed as co-constitutive modes of spectral phenomena.

Through connections and due distinctions between these two thinkers, and by drawing on insights from practice, the last part of the chapter outlines the theoretical framework I describe with the term *hauntography*. At its core, hauntography operationalises Derrida's hauntology through the pragmatic categories provided by Fisher. These inform a circle of praxis combining an analytical approach, which focuses on the affective encounter with the spectral, with a creative one, demonstrating, to paraphrase Derrida, how to give back speech to ghosts in one's work (Derrida 1994a, 221). This framework, whose applicability extends into media studies, cultural studies, science-technology studies, and potentially to other fields, constitutes the theoretical and methodological backbone for the analysis of music AI in Chapter 7, where the third research question is addressed.

6.2 A Hauntographic Glossary

As a way of scaffolding the reading, this section collects into a glossary all the vocabulary relevant to hauntography. It mainly focuses on notions developed in this chapter, but also includes relevant notions appearing throughout the dissertation. The glossary is intended as a tool for consultation rather than for sequential reading, as the framework binding all the notions together will be introduced throughout the chapter: the reader is invited to return to this section as concepts are encountered and applied.

Before proceeding, a word of caution is due. Derrida rarely offers positive definitions of the terms he develops. This is a rhetorical strategy which underpins his critique of language as presence. While the following definitions draw from hauntology theory, they should be viewed as possible operationalisations rather than tentative approximations of Derrida's own thinking.

Eerie: The *eerie* is (together with the *weird*) one of the two ways through which we apprehend or conjure the *spectre* within a certain *hauntogram*. It is characterised either by a failure of presence or a failure of absence.

Exorcism: A performative act that conjures something into presence by declaring it absent (or conjures something into absence by declaring it present). Following Derrida, exorcisms are performative attempts to declare a death only to put to death, and inevitably provoke the return of what is conjured away. Such is, for instance, Fukuyama's "end of history."

Ghost: The *ghost* embodies the vivifying principle behind the *spectre*. It represents an elusive potential that predates ontological differentiation and manifests its presence-absence through the *spectre*.

Hauntogram: The subject of hauntographic observation, encompassing a certain object or phenomenon and the context in which it is apprehended. Since *hauntography* views things and phenomena as *spectral*, the *weird* and the *eerie* are always latent in a certain *hauntogram*, and emerge as the relation between the object and the context shifts.

Hauntography: A posthumanist framework for the analysis and mobilisation of the *traces* (*spectres*) of what operates by being absent (the *ghost*) and is apprehended through the affective and fictional categories of the *weird* and the *eerie*. Thanks to this fictional and affective duality, *hauntography* can be applied in both creative and analytical modes. These modes configure a circle of praxis where practice and theory inform each other.

Presence: In hauntographic terms, *presence* is assumed as never fully present, and inherently haunted by the absence it excludes. Presence always entails absence and vice versa.

Principle of Radical Alterity: An epistemological stance whereby the *weird* and *eerie* outside is used to destabilise the inside, opening new ways of knowing.

Spectrality: A condition of ontological and epistemological liminality, or in-betweenness. Hauntologically speaking, spectrality is inherent to all ontologies, since these are haunted by the negatives that sustain them. In *hauntography*, the term describes the field of relations of ghosts and spectres, and the ways in which these are apprehended or conjured in a *hauntogram*.

Spectre: The *trace* and phenomenal manifestation of the *ghost*. It is not a subject, but rather a paradoxical disturbance of presence and absence, the experience of a “tangible intangibility” that defers to the *ghost*.

Spectrophonia: Spectrophonia describes a *weird* disjuncture of schizophonic sounds, as these are pulverised and reconfigured anew, and the *eerie* familiarity in the algorithmic outputs, haunted by the *spectres* of the multitudes they derive from (*ghosts* in the hauntographic sense, since they are in a state of continuous deferral). While anticipated here for completeness, this notion is developed in Section 7.4.

Threshold: A gate between worlds and worldviews. Thresholds separate worlds and worldviews while at the same time putting them in dialogue. In hauntography, thresholds allow the outside to enter and subvert familiar structures and thus operationalise the *principle of radical alterity*.

Weird: The *weird* is (together with the *eerie*) the affect and mode of fiction through which we apprehend or conjure the *spectre* within a certain *hauntogram*. It is characterised by a sense of wrongness, a suspension or violation of the limitations of time, space and natural law. One experiences the *weird* as a more or less temporary vertigo caused by the realisation that the categories we used to make sense of the world are now inadequate.

6.3 Hauntology

6.3.1 *Spectres at the End of History*

In the final decades of the twentieth century, as the Soviet Union crumbled and the political world order increasingly shifted in favour of the Western bloc, a singular thesis began to emerge among members of the Western intelligentsia. Francis Fukuyama's notorious 1992 bestseller *The End of History and the Last Man* announced the end-point of humankind's ideological evolution and the universalisation of Western liberal democracy as the final form of human government (Fukuyama 2006). According to Fukuyama, freed from the burden of ideology, the empirical man could now enjoy liberal democracy and the neoliberal economic model it subsumes, as the trans-historical and universal order of a just society.

Yet, like an ouroboros, Fukuyama's end of history feeds on its own tail. As has become all too apparent (if not by scholarly work alone, then through the test of time) it represents a not-quite-covert attempt to naturalise the status quo, to conceal ideology at work as background noise, to return to the age of myth and to a time beyond time. In *Specters of Marx* Derrida embarks on the timely task of deconstructing Fukuyama's thesis. Derrida writes:

This end of History is essentially a Christian eschatology. It is consonant with the current discourse of the Pope on the European community: destined to become a Christian State or super-State, this community would still belong therefore to some Holy Alliance (Derrida 1994a, 76).

Political (Holy) alliances, Derrida argues, often represent collective attempts to exorcise away some type of ghost. Fukuyama's urgency in proclaiming a time-without-time performs, indeed, an act of exorcism. This performativity is described as a form of *staging* (Derrida 1994a, 10), which in posing the question of the end conjures away the very possibility of a different beginning. In this light, Fukuyama's prophecy acquires a more sinister nuance: in eloquently witnessing the end of history, he is engaging in a murder.

The body, and the inevitable ghost arising from the conjuration, is that of Marxism. Yet crucially, communism was already a spectre in the first place, haunting Europe from the very first lines of the *Manifesto of the Communist Party* (Marx and Engels 1979), and waiting for the right conditions to haunt again. "After the end of history, the spirit comes by coming back,"

and becomes a revenant (Derrida 1994a, 10). It is precisely in this movement, the performative conjuring away of a dread (the communist spectre), which in this very act of exorcism becomes all the more present, that Derrida's notion of *hauntology* finds its political foundations.⁸²

This movement of exorcism and return of the exorcised carries with itself a question: whither Marxism? In what forms has it survived the onslaught? The occasion to investigate these questions came in April 1993, one year after the first publication of Fukuyama's *The End of History and the Last Man*, at the Centre for Ideas and Society at the University of California, Riverside. Here, Derrida gave a plenary lecture responding to the theme of the conference: *Whither Marxism? Global Crises in International Perspective*. The lecture was later expanded into the 1994 monograph *Specters of Marx: The State of the Debt, the Work of Mourning, and the New International*.

When discussing Derrida's philosophical and political contribution to the looming question of history's end, we are thus confronted with two overlapping works, that nonetheless reflect on Marx's legacy from slightly different angles.⁸³ While the foundational themes of hauntology are introduced in the lecture, the term itself is introduced in the monograph. Conversely, some themes that are paramount in the lecture (in particular, the deconstruction of the rhetoric of the end) give way in the book to a much broader philosophical engagement with spectrality as a deconstructive expedient. The sections that follow mirror this trajectory, beginning with the lecture's politically urgent arguments, and expanding them through the monograph.

6.3.2 *The Mode of Being of a Spectre*

Derrida opens with an anticipatory conjuration, a theatrical move announced right from the epigraph. The spectre of Marxism had been there all along, hidden in plain sight at the very beginning of the *Manifesto of the Communist Party* (Marx and Engels 1979), waiting for this very moment to fully reveal itself.

82 For Derrida, the act of conjuring away is ultimately one of exorcism. A performative that always causes the ghost to come back, that "convoke[s] the revenant that [it] conjures away" (Derrida 1994a, 175)

83 The New Left Review text I refer to here in discussing the lectures (Derrida 1994b) is an authorised abridgment of Derrida's oral presentation.

I had not done so for decades [reading the Manifesto]—and that must tell one something. I knew very well there was a ghost waiting there, and from the opening, from the raising of the curtain (Derrida 1994b, 32).

This theatrical incipit serves at least three purposes. First, it endows the author's deconstructive enterprise with an aura of urgency (why now? One is left wondering). Second, from the "raising of the curtain," Derrida intends to introduce, at a meta-narrative level, a performativity that while initially crucial in interpreting Fukuyama is then progressively expanded to all ontologies. Third, through this performativity, Derrida foregrounds the theme of the distorted and disjointed temporality that characterises spectral apparitions (the ghost was there all along, predating the conjuration). This latter point represents a crucial feature of spectral phenomena. Indeed, in discussing the mode of presence of a ghost, Derrida writes:

Before knowing whether one can differentiate between the spectre of the past and the spectre of the future, of the past present and the future present, one must perhaps ask oneself whether the spectrality effect does not consist in undoing this opposition, or even this dialectic, between actual, effective presence and its other (Derrida 1994b, 36).

Derrida's discourse on spectrality may be broadly viewed as a critique of the metaphysics of presence as tied to a certain way of thinking about presence and time. This critique, along with a demystification of Fukuyama, entails a distancing from Marx as well, who does not like ghosts and yet "thinks of nothing else," who conjures them only to "exorcize and analyze the spectrality of the spectre" (Derrida 1994b, 37). In other words, Marx haunts and is haunted.

There is, clarifies Derrida, more than one spectre of Marxism. "There must be" more than one of them, he suggests, implying a stringent urgency (Derrida 1994b, 33). This is indeed Derrida's working thesis and, admittedly, his bias. The question "Whither Marxism?" may thus be extended as follows: How should the West recuperate the spirit of Marxism after what seemed to be the dissolution of its political enterprise? From which of its many threads should one pull the ghost's cloth? Derrida unravels at least three threads here, addressing respectively the intellectual, the techno-mediatic, and the political dimensions.

As for the intellectual enterprise, Fukuyama is taken as the champion of a complacent intellectual establishment, with the deconstruction of his thesis on "the end of history" revealing the disjunctions of a truth presented as inevitable ("the good news *has* come," writes Derrida (Derrida 1994b, 43)).

The second instance draws a parallel between Fukuyama's rhetoric *as* staging, as "a symptomatic montage that responds to a demand, in order to reassure it," (Derrida 1994b, 49) and the spectralising powers of the techno-mediatic apparatus. To exemplify this point, which is crucial for the application and expansion of hauntology into the analysis of the AI medium, let us for a moment abandon Derrida but in order to clarify Derrida, and turn to a uniquely terrible and revealing example of media's spectralising powers provided by Giorgio Agamben. Three years before Derrida's lecture, in his notes to Guy Debord's *Comments on the Society of the Spectacle*, Agamben describes the 1989 Timisoara genocide as follows:

For the first time in the history of humankind, corpses that had just been buried or lined up on the morgue's tables were hastily exhumed and tortured in order to simulate, in front of the video cameras, the genocide that legitimized the new regime. What the entire world was watching live on television, thinking it was the real truth, was in reality the absolute nontruth; and, although the falsification appeared to be sometimes quite obvious, it was nevertheless legitimized as true by the media's world system, so that it would be clear that the true was, by now, nothing more than a moment within the necessary movement of the false. [...] Timisoara is, in this sense, the Auschwitz of the age of the spectacle: and in the same way in which it has been said that after Auschwitz it is impossible to write and think as before, after Timisoara it will be no longer possible to watch television in the same way (Agamben 2000, 80).

Agamben's notes present us with a vivid and terrifying example of the spectralising powers that the techno-mediatic apparatus begins to exhibit precisely in the historical conjuncture when Derrida writes, and of how these can serve political (and even Holy) alliances. Here, one is confronted with a many-layered performativity that reverses cause and effect (and thus, historicity) in the event and beyond the event (it does not suffice to falsify: it has to be clear, now and for the future, that the true is a moment of the false):⁸⁴ first, the exhibition of the exhumed corpses on television; second, the "symptomatic montage," to use Derrida's own words on Fukuyama (Derrida 1994b, 47), of the tortures; third, the staged legitimisation of a blatant untruth through a narrative that, much like Fukuyama's, seems to respond "to a demand" and "on demand." After Timisoara, it becomes impossible to understand media the same way, argues Agamben. It is not a stretch to imagine Derrida arguing that herein lies one reason to think of media technologies hauntologically.

84 This also echoes Jameson's discussion on nostalgia, and its displacing of the subject outside historical time.

Finally, the third and concluding thread is fundamentally a political one, aimed at recuperating Marx within the structure of social democracies. This thread would deserve a thorough analysis and critique (perhaps through the hauntological tools that Derrida himself offers), which extends well beyond the scope of this dissertation. Here, Derrida focuses on recuperating one of the spirits of Marxism not just as philosophy but as political action, through the establishment of a New International: a loose, non-homogeneous, and informal network of people and movements sharing a common concern for justice, democracy, and human rights.⁸⁵

6.3.3 *Hauntology*

In the monograph, Derrida frames a series of terms around spectres and ghosts, the first of which is the notion of *hauntology*. A term that nonetheless appears no more than three times in the 248 pages of the text.

The first passage mentioning hauntology concerns the staging of the rhetoric of the end, now abstracted from Fukuyama's critique into a critique of ontology as such. This critique is framed around the impossibility for what is discursively performed to include its own terminality *beyond* discourse, precisely because in the very act of being performed, the end is also inevitably deferred. Accounting for this paradox requires extending *being* into something that both comprehends and transcends it, that is, by expanding one's view from ontology to *hauntology*:

Staging for the end of history. Let us call it a hauntology. This logic of haunting would not be merely larger and more powerful than an ontology or a thinking of Being [...] It would harbor within itself, but like circumscribed places or particular effects, eschatology and teleology themselves. It would comprehend them, but incomprehensibly (Derrida 1994a, 10).

The second passage extends hauntology into a category of technics, now addressing modern media as inherently spectral. Importantly, media technology does not merely perform spectrality as in Agamben's passage. Rather, haunting is here viewed as inherent to the structure of the technology itself. The "medium of the media" is neither living nor dead.

⁸⁵ Here, it is inevitable to observe that class struggle, foundational to historical materialism, seems to have become itself a spectre.

Presence is constructed, reproduced, diffused and deferred into the public sphere by a dispositif that escapes ontology in its very nature.

The medium of the media themselves (news, the press, telecommunications, techno-tele-discursivity, techno-tele-iconicity, that which in general assures and determines the spacing of public space, the very possibility of the *res publica* and the phenomenality of the political), this element itself is neither living nor dead, present nor absent: it spectralizes. It does not belong to ontology, to the discourse on the Being of beings, or to the essence of life or death. It requires, then, what we call, to save time and space rather than just to make up a word, hauntology (Derrida 1994a, 63).

Finally, the third and last passage mentioning hauntology advocates the need to introduce haunting into the ontology of ideas. This move places hauntology at the centre, and ontology at the periphery, viewing the latter as the act of exorcism by which spectres are conjured *into* and *away from* presence.

To haunt does not mean to be present, and it is necessary to introduce haunting into the very construction of a concept. Of every concept, beginning with the concepts of being and time. That is what we would be calling here a hauntology. Ontology opposes it only in a movement of exorcism. Ontology is a conjuration (Derrida 1994a, 202).

It is clear how Derrida views hauntology as a necessary category to understand the present. A deconstructive category by all accounts, that similarly to terms such as *trace* and *différance* alludes to how nothing enjoys a purely positive existence. Ontologies are exorcisms insofar as they rely upon (and yet conjure away) the negatives that predate and surround them. These negatives, particularly in the case of hauntology, haunt from a place of disjointed temporality:

“The time is out of joint”: time is disarticulated, dislocated, dislodged, time is run down, on the run and run down [traqué et détraqué], deranged, both out of order and mad (Derrida 1994a, 20).

A temporal disjuncture as the trace of a haunting. Time disarticulates when its unfolding contradicts one’s understanding of empirical reality. The spectral apparition is foregrounded by the *out of joint*: the severing, the reversing, the undoing of the thread of time, the trauma of a fractured historicity. This out of jointness, and more generally, the *out of*, will prove crucial

for bridging Derrida's hauntology with Fisher's notions of the *weird* and the *eerie* (see Section 6.5.3).

6.3.4 *Three Things of the Thing*

Derrida characterises the hauntological apparition, the *Thing*, be it a spectre or a spirit "in the sense of the ghost in general" (Derrida 1994a, 6) by three features, which effectively anticipate some of the parallels and distinctions with Fisher drawn in Section 6.5:

Mourning: The attempt to "ontologize remains" to make them present (Derrida 1994a, 9), to localize them, thus making them dead once and for all by exorcising their return. What is exorcised, here, is not knowing. This brings us to the visor (and auditor) effect (see Section 1.2), to the presence that arises from what cannot be fully known, and the gaze that returns ours from the position of asymmetry of what sees without being seen.

If the spectre is conjured through an act of mourning, mourning in turn "always follows a trauma" (Derrida 1994a, 121). Trauma is signalled by and signals a temporal disjunction. As discussed in Section 6.4.2, Fisher draws extensively on trauma as a fracturing force in his framing of *sonic hauntology*, and further in his parallel between hauntology and Afrofuturism (Fisher 2013b).

Language: Crucially, the *Thing* is an act of language; it is an absence made present through a signifier, that "marks the name or takes its place" (Derrida 1994a, 9). Signification, or at least the condition of its possibility, is also the condition for the ghost's promissory function, whose apparition introduces in and of itself the need for a reckoning. This spectralising power of language (and of AI-generated sounds taking on its functions) emerged with particular strength as I worked on *Mouja+* (see Section 5.4).

Work: The ghost manifests itself as *work*, in its ability to transform itself and the world around it. This disembodied power of transformation foregrounds another convergence between Derrida and Fisher. As Fisher aptly summarises, hauntology is about "the agency of the virtual," with the spectre understood as that which acts without (physically) existing (Fisher 2013a, 27). This echoes Fisher's own framing of the *eerie* as either a failure of presence or a failure of absence, as further discussed in Section 6.5.2 (Fisher 2016, 61).

To conclude this necessary overview of relevant hauntological notions, one last terminological clarification is due. It is time to ask ourselves, for lack of a better term, how to name the *Thing*, the phantasmatic operator, the subject of hauntological deconstruction. While Derrida never definitively resolves the ambiguity between spectres, spirits and ghosts, the most clear-cut distinction he provides is the one between the spectre and the spirit-ghost:

What distinguishes the specter or the revenant from the spirit, including the spirit in the sense of the ghost in general, is doubtless a supernatural and paradoxical phenomenality, the furtive and ungraspable visibility of the invisible, or an invisibility of a visible X. [...] It is also, no doubt, the tangible intangibility of a proper body without flesh, but still the body of someone as someone other (Derrida 1994a, 6).

The spectre “begins by coming back.” It is not a subject but a paradoxical disturbance, the “tangible intangibility” of a deferred “other.” The ghost (or the spirit that Derrida describes as a specific instance of the ghost) embodies the vivifying principle that precedes and produces the spectre. In other words, while the ghost is an elusive potential that predates any ontological differentiation (in a way, a deferred principle haunting *différance*), the spectre is its paradoxical incorporation, the becoming-body, a certain “phenomenal and carnal form of the spirit,” and thus its trace in the world.

6.4 Towards Sonic Hauntology

In a closing passage of the 1989 essay, Fukuyama extends his argument around the “end of history” into the “end of art,” arguing that, as humanity enters a time without history, art and philosophy would inevitably have to resign themselves to the “perpetual caretaking of the museum of human history” (Fukuyama 1989). Fukuyama’s thesis draws on Kojève’s reading of Hegel, paraphrasing the former to suggest that, as history comes to a close, “there would be no new eras and no particular distinction of the human spirit for artists to portray” (Fukuyama 2006, 311). This indirect characterisation arguably distorts Hegel’s thinking. For the German idealist, there exists no endpoint to artistic creation and no shift in the relevance of artistic production. The shifting function of the arts is rather understood within a mechanism of

overcoming, as inherent to a dialectical progression. Here, art cannot end, but changes precisely so as not to lose its relevance amidst mutating conditions.⁸⁶

A certain sense of terminality is also inseparable from the cultural current embodying the end of the grand narratives that goes by the name of *postmodernity*. Lyotard describes postmodernity as an aesthetic, cultural and political movement, a non-ideology characterised by a fundamental scepticism towards past narratives and a generalised loss in the stability of meaning (Lyotard 1984). This condition turns the artist into a philosopher, as the work of art, in the absence of a tradition to refer to, defines its own methods and categories.

While recognising the tendency of postmodern art to imperceptibly turn into its own theory and the theory of itself, Jameson underlines a fundamental paradox in postmodernity's framing as a non-ideological and a-historical discontinuity:

The representational paradoxes involved in any narrative whose fundamental category is the postmodern "group" can then be articulated as follows: since the ideology of groups comes into being simultaneously with the well-known "death of the subject" (of which it is simply an alternate version) [...] the consequence will be that these new collective characters and representations that are groups cannot any longer, by definition, be subjects (Jameson 1991, 348).

According to Jameson, postmodernism's claims are evidently burdened by the weight of their own logic of dissolution. This contradiction places the subject outside of history, and becomes evident in what Jameson terms the *nostalgia mode*, a longing for an aestheticised past escaping historicity. This reconstructed, stereotyped and re-actualised version of the past is achieved through the recovery of its simulacra, and "endows present reality and the openness of present history with the spell and distance of a glossy mirage" (Jameson 1991, 20). Nostalgia does not merely replace a real temporality with a fictitious one; it rather foregrounds a crisis in the representation of the individual's own historical situatedness. As memory becomes a repository of disconnected images, styles and affects, the present becomes the market in which a disjointed self, of which Gumm in Dick's novel is a quintessential example, consumes them.

While Lyotard celebrates postmodernism as a non-ideological resolution of the modernist contradictions, Jameson insists on its ideological continuity with that of the "postindustrial society." Jameson writes:

⁸⁶ According to Hegel, the historical function of sensuous art "has rather been transferred into our ideas instead of maintaining its earlier necessity in reality and occupying its higher place" (Hegel 2010, 11).

Lyotard's aesthetic positions, however, cannot be adequately evaluated in aesthetic terms, since what informs them is an essentially social and political conception of a new social system beyond classical capitalism (our old friend "postindustrial society"): the vision of a regenerated modernism is, in that sense, inseparable from a certain prophetic faith in the possibilities and promise of the new society itself in full emergence" (Jameson 1991, 59).

In conclusion, by extending Jameson's own words on postmodernism, the "end of history" and postmodernism's "end of grand narratives" together represent the "reflex and the concomitant of yet another systemic modification of capitalism itself" (Jameson 1991, 59) that displaces the subject into a timeless (and thus, out of joint) present.

6.4.1 *Capitalist Realism and Hauntological Resistance*

Having analysed Derrida's hauntology and outlined how the notion of the end of history infiltrates culture and the arts, we are now in the best position to introduce Fisher's discourse on *sonic hauntology*, which applies Derrida's thinking to music.

The political context in which Fisher writes follows the one Derrida engages with. Fisher lives in and describes the post-Thatcherite UK, in the aftermath of the Western bloc's alleged victory against the communist spectre, a time (our time) when history itself has proven Fukuyama's claims unmistakably false.

Yet prophecies, especially when sustained by capitalism's notorious ability to assimilate its own critique, find peculiar ways to fulfil their promises. As practices and rituals are progressively aestheticised and commodified, and as the creative potentials of the phantasmal world are increasingly conquered by the Unnamable Thing (Deleuze and Guattari 1983, 154), it becomes "easier to imagine the end of the world than it is to imagine the end of capitalism" (Fisher 2010).

The discourse of the end transposes here from historical consciousness, against which it cannot empirically stand, to the unconscious, where it can exist independently of any reality check and work its way into fulfilling its promise. From this particular angle, however rightfully debunked, Fukuyama's thesis is still enacting the prophecy disclosed by Derrida, not because of persuasive power or due to historical timeliness, but due to a much broader process of ideological naturalisation through which the economic model has installed itself in people's

“hearts and souls,” as Thatcher famously stated, operating as a libidinal force beyond time, empirical reality, and history. It is thus not by “making history right” that one confronts this spectre, but by re-appropriating desire through the building of a new imaginary, one that subverts the current hegemony through its very means. This is the move of *sonic hauntology*.

6.4.2 *Sonic Hauntology and the Metaphysics of Crackle*

To exemplify how cultural time has folded onto itself, Fisher proposes a thought experiment. Imagine a twenty-first-century music recording being beamed back to 1995, and ask yourself whether its newness would shock the listeners as much as a jungle recording being sent back from 1993 to 1989. The shocking novelty of the latter, as opposed to the appalling aesthetic conformity of the former, is, according to Fisher, symptomatic of a closure into the nostalgic repetition of the known, and the coming of a future that “doesn’t feel like the future” (Fisher 2013a, 18).⁸⁷

In the current state of affairs, where the cultural imaginary forecloses the construction of desirable futures beyond neoliberalism, and media technologies collapse space and time into a disjointed present, Fisher’s *Metaphysics of Crackle* addresses Derrida’s discourse on presence from the producer’s desk. In explicit continuation of Derrida’s philosophical project, the “crackles” Fisher refers to are the hiss, pops, and noises of the phonographic support, traces artificially reintroduced and layered onto the transparency of the modern digital medium by a group of artists described as *sonic hauntologists* (Fisher 2013b).

In the works of William Basinski, The Caretaker and the artists gravitating around the British *Ghost Box* label, these artefacts unsettle the very distinction between surface and depth, between background and foreground, thus instantiating a series of unsettling temporal disjunctions. According to Fisher, this artifice operates in two ways: first, in the temporal domain, by signalling the disjunction between the modern digital medium and the phonographic one, thus alerting the listener of the medium’s presence; second, in the ontological domain, by introducing the technical frame as part of the content (Fisher 2013b).

87 While I agree with Fisher’s conclusions, I approach this thought experiment with some caution. Several objections could be advanced here. My main objection is relative to the locality of the claim. Interestingly, Fisher’s own example looks at jungle music, a genre that entered the UK from Jamaican roots. Looking beyond the imperial core, the cultural exhaustion described by Fisher gives way to a much more lively experimentation.

Sonic hauntologists explicitly address the question of memory after the advent of the ideology-in-disguise that is the end of history, reacting to the nostalgia mode by deconstructing the cultural imaginary into a myriad of impossibly contiguous “nows” that reveal the ideological palimpsest criticised by Jameson. As Fisher puts it, “whereas [nostalgia] glosses over the temporal disjunctures, the hauntological artists foreground them” (Fisher 2013b). This is achieved through an extended use of sampling techniques, reconfiguring electronic sounds from the second half of the twentieth century, in themselves signifiers of the future, into dreamlike, eerie soundscapes. Through this artifice, nostalgia is subject to a paradoxical reversal, where the longing is displaced ahead, into the futures that never came to be as a consequence of postmodernity’s temporal closure.

Crucially, sonic hauntology transforms hauntology’s deconstructive approach into an active instrument for artistic creation and political dissent. Temporal disjunctions, here as much as in Derrida, foreground trauma (see Section 6.3.4). Yet, in Fisher trauma represents much more than the canary in the coal mine revealing ideology at work, or the opening for an indefinite and unknowable potential of change from within the unassimilable (what Derrida describes as a messianic without messianism). Rather, trauma represents the occasion to override the present’s cultural closure and displace one’s critique from the future (a move that characterises both sonic hauntology and Afrofuturism), the catalyst for a resistance through culture that begins *ahead* of the disjunctions of a time without time.

6.5 The Weird and the Eerie

In *The Weird and the Eerie*, Fisher develops what appears to be an entirely separate project, distinct from his writings on sonic hauntology, and even more so from Derrida’s hauntology. The intent of this work is to analyse what Fisher describes as two categories of the *strange*, namely the *weird* and the *eerie*, which represent modes of artistic production as well as affects that one encounters in culture and “in the raw” of natural phenomena (Fisher 2016).⁸⁸ Yet, in reading *The Weird and the Eerie* in parallel with *Specters of Marx*, it becomes evident how these two categories (along with many of the sub-categories that Fisher identifies) share

⁸⁸ This distinction is functional rather than ontological. The aim here is to show how one can use the *weird* and the *eerie* creatively, or simply register them all around, not to advance claims on nature-nurture dualisms.

hauntology's fascination with what haunts presence, with the gaps in our understanding of reality, with the "outside" and the "beyond."

In this section, I advance the following arguments: First, that *The Weird and the Eerie* may be effectively understood in continuity with Fisher's discourse on sonic hauntology; second, that the *weird* and the *erie* exhibit noteworthy convergences with Derrida's framing of ghosts and spectres; third, that the *weird* and the *erie* may be used to operationalise hauntology into a framework to engage with phenomena through the lens of spectrality. By following a series of textual cues from Fisher's broader corpus and *Specters of Marx*, the following sections establish these two modes as part of the theoretical foundation supporting the hauntographic framework.

6.5.1 *The Weird*

The *weird*, according to Fisher, is apprehended as a sense of wrongness, of suspension, or violation of the "galling limitations of time, space and natural law" (Fisher 2016, 18). One experiences the *weird* as a more or less temporary vertigo caused by the realisation that the categories we use to make sense of the world are inadequate. This typically happens when something "other" or "alien" infiltrates ordinary experience, cracking the solidity of what is taken as given.

David Lynch's *Mulholland Drive* (Figure 6.1) offers compelling examples of the *weird*, both as a mode of artistic fiction and of perception. In one of the most iconic scenes, the protagonists Betty and Rita enter *Club Silencio*, where a magician-compere takes the stage with the grave announcement: "No hay banda!" (There's no band!). The viewer knows early on that what they are witnessing is pure fiction, that all sounds are recorded and played back. Yet, it is impossible not to shiver as the trumpet player theatrically raises the instrument above his head while a sustained note continues to resonate through the space. The deception, albeit theatrically announced, proves inescapable a second time, when the singer Rebekah Del Rio collapses on the stage while her passionate singing persists.



Figure 6.1. *Club Silencio, Mulholland Drive (2001).*

What makes this scene *weird* rather than merely surprising is how it plays with the crossing of *thresholds*, foregrounding multiple *beyonds*, other realities that infiltrate and undermine the solidity of the one we inhabit as viewers: first, the thresholds between fiction and reality represented by the club's entrance, the stage, and Lynch's signature red curtain; second, the sounds in themselves as thresholds to the "beyonds" of the sources they were severed from and which they schizophrenically seek to rejoin through new embodiments; third, a threshold within cognition itself, in that despite being repeatedly told and shown that the sounds we hear do not belong to the bodies on the stage, we are nevertheless compelled to reconcile them with what is available to our apprehension. This latter point is of particular interest to the *weird*, in that it reframes the human from the perspective of the processes of inert matter, by revealing that cognition is never entirely self-owned, but is governed by forces that transcend the subject (Fisher 2016).⁸⁹

Importantly, this scene also introduces a form of montage, which Fisher describes as a juxtaposition of two things that do not belong together (the sounds and their new

⁸⁹ A realisation that Fisher parallels to the encounter with the Lacanian Real.

embodiments), and that, he argues, represents the form that is most appropriate to the *weird*. This consideration reveals an evident continuity between the *weird* and Fisher's earlier reflections on sonic hauntology. As we saw, the methods of political critique upon which sonic hauntologists ground their practice are founded precisely on reconfiguring sonic fragments from distant temporal planes to foreground the possibility of a future in a time without history. In other words, they rest on a form of montage.

If montage offers a bridge between the *weird* and Fisher's hauntological reflections, a clue for connecting the *weird* to Derrida's hauntological categories may be found in the phrase *out of*, which is crucial for both thinkers. Indeed, as we saw in Section 6.1, Derrida opens *Specters of Marx* by citing Hamlet's phrase "time is out of joint," and the same phrase returns both in Fisher's discussion of Philip K. Dick's novel (Dick 1984) and in his characterisation of Lovecraft's writing.

To understand the *out of*, it is important to stress the consideration advanced above, that in the *weird* the "here" and "there" always engage in a dialogue. The *out of* embodies this dialectic, in that it may be interpreted at the same time as "from," as "cut out of," and as "beyond" (Fisher 2016, 11). In *Specters of Marx*, the apparition of Hamlet's ghost renders time *out of joint*, in the sense of "disarticulated, dislocated, dislodged, deranged, both out of order and mad," with the ghost as the agent of a weird subversion of empirical reality, which Derrida parallels to the political disjunctions, the traumas of our time. In other words, for Derrida it is through the *out of joint* that the deferring principle of the ghost is apprehended in the world.

Derrida invites the scholar of the future to engage with the spectres that materialise the ghost's vivifying force, to give them speech and receive it back. Fisher shows how these manifest in one's experience, as *weird* (and *eerie*, as discussed hereafter) phenomena, and ventures further by explaining how to conjure and analyse this affect through a series of sub-categories, a non-exhaustive list of its different variations within contemporary art and literature. Among such are: the Fall's grotesque style, that reframes mainstream rock aesthetics through the parody of a "psychotic outside;" Dick's un-worldings, where truth infiltrates a false reality from within; the "real externality" of Lovecraft's "out of space and time;" the ouroboros effect of time folding onto itself in Tim Powers.⁹⁰

In all of the above cases, the common element of novelty that the *weird* introduces lies in how it foregrounds an encounter with the outside, not by colonising it, but rather by letting

90 For an in-depth analysis of these sub-categories of the *weird* refer to Fisher's monograph.

it enter our reality through thresholds of different kinds. But to become spectral, at least in the hauntographic framing I advance here, the weird needs to encounter the *eerie*.

6.5.2 *The Eerie*

As in the case of the *weird*, the *eerie* can be either conjured or encountered in the raw. This can happen with artistic and technical artefacts such as Thales and Stacco (see Chapter 4), in performances such as those described in Chapter 5 (it was particularly evident in *Mouja*, as discussed in Section 5.2.4), but also in a natural environment almost untouched by human hand. As a counterpoint to the example of the *weird* offered above, the example I use to introduce the *eerie* falls into the latter group. Specifically, I draw on the Bosque Quemado (Figure 6.2), an inhospitable area on the slopes of the Turrialba volcano in Costa Rica, which I was fortunate to visit as I was writing these pages.

At more than 3000 metres above sea level, this unique land is shaped by an agent hiding behind a misty background, tracing its presence through a black soil and skeletal trees punctuating a strikingly colourful, receding vegetation. In walking the trails short of breath due to the height, one is initially struck by the unusual contrast between the vibrant green and the sparse trees, standing as dark omens in an otherwise bewitching land. By this description, it is clear how the estrangement that welcomes the visitor involves some element of the *weird*. Yet, this sense of wrongness foregrounds the presence, concealed behind the mist, of a volcano too extended both spatially and temporally to reveal itself in its fullness.⁹¹

91 While hyperobjects (Morton 2013) are arguably eerie, it is relevant to note that not all eerie phenomena are hyperobjects.



Figure 6.2. *Bosque Quemado, Turrialba volcano, Costa Rica.*

The gaps in the receding forest bear witness to an elemental friction within nature, a tension whose spatial and temporal scales can hardly be comprehended by the occasional excursionist. In venturing into the area, one cannot avoid wondering what beings would populate this place, what minuscule presences one would find through the cracks in the burned wood, or what fantastic gardeners might be tending to the uncannily even hedges on the trails' sides.

If the *weird* may be described as the affect of that which does not belong and alters our understanding of the ordinary, the *eerie* may be described as a failure of presence or a failure of absence. It raises the fundamental question of who, or what, is the entity we perceive (Fisher 2016, 12). This question inevitably echoes Walshe's question of who is listening and how through the hidden layers of an AI model (Walshe 2023).

Thus, the *eerie* links presence with knowledge (and brings us back full circle to the overarching narrative of Chapters 4 and 5), so much so that when enough understanding is achieved (I think for instance of the farmers cultivating potatoes all around the slopes, knowing all too well the agents that shape the land), the *eerie* disappears. This does not mean that it dissipates once and for all, though. It rather recedes into a dormant state, ready to be conjured

through shifts in context that reveal the phenomenon anew. The volcano's distant rumbling after months of silence, for instance, would send shivers down the spine of those working the land, suddenly reminding them of its latent potential.

As with the *weird*, Fisher draws examples of the *eerie* from films and literature to introduce a series of relevant subcategories. These include: confabulation, where the *eerie* emerges from the agency of the unconscious, from “the conundrum of how it is possible to conceal something from ourselves” (Fisher 2016, 72); the sublime provoked by glimpses of forces too large to be fully apprehended; acousmatic sounds, in which voices or noises persist without their source being present (in this sense, the Club Silencio scene combines the *weird* and the *eerie* into a fully spectral manifestation); negative hallucinations, in which the sensory apparatus fails to experience something that is present, resulting in experiential gaps; alien traces, suggesting agencies that operate according to non-human logics; the dissolution of subjectivity through what Fisher calls the “outside-in” of an alien perspective infiltrating human consciousness, or, vice versa, the “inside-out” of an estranged human subjectivity.⁹²

In all these manifestations, it is clear that the *eerie* ultimately raises a problem around agency and presence, which, in both Fisher and Derrida, do not coincide. As Derrida observes in his characterisation of the *Thing* (see Section 6.3.4), the ghost exhibits “a certain power of transformation.” It is *work* without presence. Similarly, the Fisherian *eerie* points to the invisible forces, human and non-human, that shape our lives from the outside-in but also from the inside-out (see Fisher's discussion on Thanatos in Kneale's *Quatermass* series (Fisher 2016, 83)), and whose agency is apprehended only through the traces they leave as they shape our world.

According to Fisher, the *eerie* becomes a crucial construct in a “tele-connected” capitalist world where agency distributes with unprecedented fluidity and scale. This further parallels Derrida's hauntological considerations on the “medium of the media,” whereby media technologies increasingly necessitate a thinking that supersedes presence and ontology, and thus requires hauntology. As will become apparent in Chapter 7, generative technologies introduce yet another turn of this *eerie* screw.

92 For an in-depth analysis of these sub-categories of the *eerie* refer to Fisher's monograph.

6.5.3 *Convergences and Divergences*

Thus far, we have analysed the *weird* and the *eerie* separately in relation to hauntology. This section moves beyond the parallels we already traced (between the *weird* and Derrida's *out of joint*, and between the *eerie* and the *Thing* as work without presence) to examine how these modes taken together converge and diverge across both thinkers' work. Here, the aim is twofold: first, to show that the *weird* and the *eerie* can coexist in spectral phenomena; second, to navigate the overlaps and differences between Fisher's and Derrida's work productively, in order to develop the hauntographic framework.

From Fisher's angle, what the *weird* and the *eerie* have in common is a certain preoccupation with the *strange* and their being at the same time modes of fiction, perception, and being. Yet, these two modes converge in a brief but revealing paragraph dealing with the notion of *fate*. Here Fisher conjures, à la Derrida, yet another of Shakespeare's ghosts:

Yet fate might be said to belong to the weird as well as the eerie. The soothsaying witches in Macbeth, after all, are known as the Weird Sisters, and one of the archaic meanings of "weird" is "fate". The concept of fate is weird in that it implies twisted forms of time and causality that are alien to ordinary perception, but it is also eerie in that it raises questions about agency: who or what is the entity that has woven fate? (Fisher 2016, 12)

In reading this paragraph comparatively with *Specters of Marx*, it is impossible not to register a certain confluence between Fisher's discussion around fate and Derrida's critique of messianism. For Derrida, messianism (that of Fukuyama and of the Holy Alliance he represents, but also of Marx) frames the future as inevitable precisely to determine it, thus implying a twisted form of causality in all respects similar to the one Fisher attributes to fate. It follows that both messianism and fate are temporally disjointed, "deranged, both out of order and mad" (Derrida 1994a, 20): in Fisher's words, *weird*. At the same time, for both thinkers, fate and messianism share a fundamental eeriness, in that they obscure the agents behind their weavings precisely to appear as natural and inevitable (in Derrida, the eschatology of the end of history discussed in Section 6.3.1).

Despite this crucial convergence, the reactions to fate's disjunctions are as different as they can be. Derrida's hauntology translates politically into an open promise and indeterminate question of justice and democracy-to-come, a "messianic without messianism" that transcends any precise expectation, guiding one "like the blind" (Derrida 1994a, 82). The vagueness of

this promise arguably explains Fisher's frustration with Derrida's epistemological posture, which he criticises for having "made a lawyerly virtue of avoiding any definitive claim" (Fisher 2013a, 25). Rather, Fisher uses the *weird* and the *eerie* to seek new forms of resistance, and as thresholds to encounter the "new."

Fisher's distancing from Derrida's deconstructive project is grounded, among other things, on a different interpretation of trauma and the return of the unassimilable. As seen in Section 6.3.4, trauma involves an element of repetition. In Derrida, the ghost's coming back, and even its "beginning by coming back," is an inescapable attractor, one that "is endlessly denied by the very movement through which one tries to cushion it, assimilate it, to interiorise and incorporate it" (Derrida 1994a, 122). On the other hand, Fisher sees trauma as a productive locus, and its unassimilability as a generative force. Fisher's parallel between sonic hauntology and Afrofuturism exemplifies this. Here, the thinker suggests that on the verge of the twenty-first century's neoliberal experiment, the unassimilable temporal disjunctures that constituted the trauma of the Afrodiasporic experience can no longer be escaped by white culture (Fisher 2013b). Time was always-already out of joint for the slave, irremediably fractured on an incommensurable scale. What remains are spectres in the form of "gaps, lost names, and screen memories" (Fisher 2013b). Similarly, time is now out of joint within the neoliberal core as well, where a socially and ecologically unsustainable present installs as the sole alternative to itself. For Fisher, Afrofuturism and sonic hauntology perform the same critique and through similar means, by displacing the discourse *from* the future (for instance, think of Sun Ra's *Space is the Place*, and of Tricky's post-apocalyptic "dysfunktionen" in *Aftermath*), and thus from outside the constraints of trauma's endless repetition.

Both Derrida and Fisher also engage with Freud's essay on the Unheimlich (Freud 1919), an engagement that in itself represents another connecting thread, yet one that the two approach from very different angles. The former criticises Freud's translation of the Unheimlich into the "terrible or the frightful" (Derrida 1994a, 218), claiming instead the ghost as its strongest example. Fisher frames the *weird* and the *eerie* as diametrically opposed to the Unheimlich. What distinguishes these two modes from it is their opposed epistemological directionality. The Unheimlich is fundamentally geared toward the strange within the familiar, subsuming an interpretation of the outside as a function of the psyche's known structures. On the contrary, paralleling sonic hauntology's transformation of trauma into a productive force, the *weird* and the *eerie* perform the opposite move, projecting the irreconcilable outside into the inside precisely to destabilise the very structures that the Unheimlich upholds.

This orientation inevitably leads into posthumanist territories, a dimension that Derrida alludes to but does not develop as much as Fisher. For Derrida, the ghost blurs the boundaries between living and inert matter (for instance, the paradox of “automatic autonomy” inherent to technical life discussed in Section 1.2). While recognising modernity’s lesson that traditional distinctions between the living and the dead do not hold any more, Fisher distances himself from the posthumanist perspective offered by the so-called new materialisms. The temptation of anthropomorphising the alien, in the sense of endowing the non-human outside with human agency, is avoided in favour of what may be described as a principle of *radical alterity*, which takes on the perspective of the outside as the default mode.⁹³ The effort, here, becomes that of exposing the invisible, transpersonal, and alien forces that govern our world, and through this turn the organic into a folding of the inorganic. In other words, in a *weird* and *eerie* posthumanism, we are mere ghosts to ourselves.

6.6 Hauntography

The above sections laid out the foundational concepts behind Derrida’s hauntology, and examined both the continuities and discontinuities with Fisher’s early writings on sonic hauntology and his later work on the *weird* and the *eerie* as modes of fiction and affects. Of particular relevance is the weaving together of the *weird*’s temporal and causal gaps with the Derridean *out of joint*, and of the *eerie*’s failure of presence and absence with hauntology’s attention to deferred and deferring presence. We are now ready to introduce hauntography, which synthesises these diverse notions into a framework for analysis and artistic creation.

I define hauntography as a *posthumanist framework for the analysis and mobilisation of the traces (the spectres) of what operates by being absent (the ghost) and is apprehended through the weird and the eerie*. While by this definition the novelty of hauntography lies principally in combining Derrida’s hauntology with the Fisherian *weird* and *eerie*, this framework incorporates and discards specific contributions from each of these thinkers.

⁹³ To further strengthen the connection between haunting, the *weird* and *eerie*, and the principle of radical alterity, it might be relevant to report this passage from Lovecraft, which Fisher cites: “when we cross the line to the boundless and hideous unknown — the shadow-haunted Outside — we must remember to leave our humanity and terrestrialism at the threshold” (Fisher 2016, 20).

6.6.1 *Heritage*

Hauntography inherits from Derrida the distinction between the ghost, as an elusive potential that predates any ontological differentiation, and the spectre, as the trace or phenomenal manifestation of the ghost. Maintaining this openness requires a certain patience, one in which exorc-analysis is resisted by letting ghosts speak and by “giving them back speech” (Derrida 1994a, 221) as a crucial epistemological positioning.

To make this argument less abstract, let us zoom out from the particular paragraph and chapter we are in, and look at this whole dissertation hauntographically. As my research on XAI progressed, the initial commitment to resolving its explanatory issues gradually shifted towards the analysis of AI’s resistance to explanation. By letting the ghosts-in-the-machine speak for themselves and inhabiting rather than conjuring them away, I could develop the artistic works discussed in Chapter 5, the notion of the haunted work, the framework discussed in this chapter, and the insights offered in the next chapter.

Equally relevant is Derrida’s attention to the inherent instability of all ontologies, the recognition that all claims of presence or absence carry the seeds of their undoing. As Derrida argues, “ontology is a conjuration” (Derrida 1994a, 202). Here, with the term *conjuring* Derrida implies the paradoxical movement of conjuring into presence to conjure away from it, and thus a form of exorcism. This view positions haunting as the precondition of all presence, and ontologies in themselves as acts of exorcism. If ghosts are always there, “even if they are no longer, even if they are not yet,” (Derrida 1994a, 221), then the question is not whether something is haunted, but what are the conditions (the contradictions, the glitches, the algorithmic “slips of the tongue” as discussed in Section 2.5.4) that make haunting apprehensible. The subject of hauntography, the *hauntogram* from now on, thus necessitates observing a certain object or phenomenon (Stacco, *End of Empire*, AI, a movie scene, a memory, a cup) in a situated manner rather than abstractly, and from there conjuring the *weird* and the *eerie* from latency (see the operationalisations in Section 7.2).

Hauntography adopts Fisher’s critique of Derrida’s tendency to recede into indeterminacy, not with the aim of superseding indeterminacy, but rather of analysing it effectively and using it creatively. Hauntology’s deconstructive foldings (the tendency to recede into a “messianic without content” that guides us “like the blind” (Derrida 1994a, 82)) are here compensated by an assertive movement grounded in what I described as a principle of radical alterity, the Fisherian orientation of adopting the perspective of the outside to reframe the inside (see Section 6.5.3).

While hauntography embraces Fisher's tension towards the new as a productive locus, a posthumanist posture that is critical of certain posthumanist tendencies to anthropomorphise the inert (see Section 7.2.3 for an example of how this translates analytically), and the categories of the *weird* and the *eerie* as affective and fictional modes, it nonetheless distances itself from Fisher in two ways: first, through Derrida, the *weird* and the *eerie* are viewed as concurrent manifestations of the spectral rather than of the *strange*; second, the *weird* and the *eerie* characterise *together* spectral phenomena, one embodying the Derridean *out of joint* and the other the inherent liminality of all ontologies.

That the *weird* and the *eerie* sow the seeds of one another is evident in Fisher's monograph, and became apparent in my practice as well. Indeed, Fisher repeatedly (even though not always) recognises weirdness in subjects he analyses as eerie, and vice versa. This is the case with the notion of fate as discussed in Section 6.5.3, of Joan Lindsay's novel *Picnic at Hanging Rock* (Fisher 2016, 122), and of Tim Powers' *The Anubis Gates* (Fisher 2016, 48). As for my practice, one example comes from how the *weird* of the gibberish produced by the human voices in *Mouja* and *End of Empire* reinforced the *eerie* effect of an unknown presence. Similarly, in *Mouja+*, the *eerie* effect caused by the physical absence of the agents conjured onstage was inevitably accompanied by a *weird* temporal and causal disjuncture: the paradox of representing something that never was in the first place (see Section 5.4.7).

6.6.2 Hauntographic Praxis

The dual fictional and affective modalities of the *weird* and the *eerie* allow one to approach hauntography from two distinct angles, which I describe respectively as *creative* and *analytical*. In a creative hauntography, one begins by working a certain hauntogram through practice, and from there conjures the *weird* and the *eerie*. Analytical hauntography begins instead from the opposite direction, from a particular observation of the phenomenal encounter with the *weird* and the *eerie*. The *creative* and *analytical* modalities identify a circle of praxis in which hauntological knowledge is gained either through a critical analysis of a phenomenon or from field work, and from there circled back into practice. While in Section 3.2 I clarified why I view hauntography as a theoretical framework rather than a method, it can be a useful exercise to detail a few possible approaches to creative hauntography and to analytical hauntography.

6.6.3 Creative Hauntography

A creative hauntography would focus on letting the generative potential of practice, whether artistic, design-oriented, or otherwise, drive one's inquiry. As we saw, ghosts "are always there, even if they are no longer, even if they are not yet" (Derrida 1994a). As practice unfolds, one remains attuned to the *weird* and *eerie* elicited by shifts between subject and context (the *hauntogram*). Once a seed of weirdness emerges, one looks into it for the *eerie*; once a seed of eeriness emerges, one looks into it for the *weird*. Once these two effects are conjured, one magnifies them through practice.

While hauntography emerges from rather than constitutes the framework for the performances discussed in Chapter 5, *Mouja+*, being the most recent piece I composed, incorporates and demonstrates part of my hauntographic thinking. In this piece, the *eerie* and the *weird* are deliberately provoked and magnified in different ways. The *eerie* is provoked by using scores (Fluxus action notation) originally intended for human actors, and by displacing sounds dynamically as if the disembodied presences they refer to were "walking" among the audience. The *weird* is achieved from here by emphasising, in a move that echoes the Club Silencio scene in *Mulholland Drive* (see Section 6.5.1), the ontological paradox of a sound (and text through the scores projected for the audience) that points to a source that did not come before it.

To offer an example beyond my own practice and in a different context, an object of design that exemplifies hauntographic thinking is Marvin Minsky's *Useless Machine*, a device whose sole function is that of switching itself off by pressing the same power switch the user presses to turn it on (Figure 6.3). As Arthur C. Clarke writes referring to the *Useless Machine*, "there is something unspeakably sinister about a machine that does nothing — absolutely nothing — except switch itself off" (Clarke 1958, 167).

Here, Clarke describes the specific case, or *hauntogram*, of witnessing the machine in use for the first time. As we saw, the *eerie* emerges when we witness something where there should be nothing, but also when we witness nothing where there should be something. The *Useless Machine* effectively embodies both these types of eeriness. The former type (something where there should be nothing) is evident in the machine's act of turning itself off right after having been turned on, upon which the user is compelled to project a certain autonomy, a "will," even a desire for not being, and thus a presence. The latter type (nothing where there should be something) depends on the device's challenge to the assumption that machines are built to fulfil a practical purpose, which here lacks altogether.

This brings us to the *weird*, which lies in the twisted causality of the *Useless Machine*'s workings. As with fate (see Section 6.5.3), we witness a destabilising, out of joint folding of causes into effects, a short circuit wherein the act of activating the machine is the very one that causes its deactivation. From here, we can even begin to observe the *Useless Machine* through the principle of radical alterity, and argue that the most frightening aspect of this device lies in the fact that its complete lack, and even refusal, of a purpose ultimately challenges assumptions about human purpose.



Figure 6.3. *Useless Machine*.

6.6.4 Analytical Hauntography

An analytical approach to hauntography invites one to observe rather than to build, and to develop a practice around the emergence of the *weird* and *eerie* in a certain hauntogram. In other words, one does not deliberately conjure a ghost, but rather looks for its spectres where they manifest. The *weird* and the *eerie* are analysed separately within a certain hauntogram, and then jointly. Finally, one may apply the additional categories of analysis identified in this

chapter (e.g. threshold, principle of radical alterity, exorcism, all of which I operationalise in the next chapter) and refine the analysis through the sub-categories of the *weird* and the *eerie*.

In Section 6.5, I listed most of the different sub-categories of the *weird* and the *eerie* as identified by Fisher, but it is crucial to note that not all of them apply directly to hauntographic analysis. This is because the latter views the *weird* and the *eerie* as co-constitutive, whereas Fisher analyses them separately. Developing a hauntographic taxonomy thus represents a substantial area for further work, a project that begins with the next chapter. While in Section 6.6.3 I offered a few examples of a creative hauntography, Chapter 7 operationalises hauntography from an analytical angle.

6.7 Conclusions

The comparative reading of Fisher and Derrida offered in this chapter aimed at formalising the theoretical framework of hauntography, which answers my second research question, concerning how the discourse and practices emerging from AI's opacity, and identified in Chapters 4 and 5, contribute to developing tools for interrogating shifting notions of agency, authorship, and intentionality in human-AI co-creation. By combining insights from these two authors, hauntography offers a practical toolset to describe and create through spectral phenomena while preserving hauntology's unique ontological posture.

I consider hauntography the core theoretical contribution stemming from this research, and define it as a *posthumanist framework for the analysis and mobilisation of the traces (the spectres) of what operates by being absent (the ghost) and is apprehended through the affective and fictional categories of the weird and the eerie*.

From Derrida's work on hauntology, hauntography inherits (i) the commitment to inhabiting the space of ontological liminality from which the *ghost* haunts, (ii) the description of the *spectre* as a phenomenal trace of the ghost, (iii) the characterisation of the *ghost* as an unknowable and receding vivifying force before differentiation of presence and absence, and (iv) a micropolitical view on all ontologies as conjurations (notably, conjuring is framed by Derrida as the double movement of conjuring into presence and conjuring away from it, and thus as an act of *exorcism*, another relevant hauntographic notion) and thus as always-already haunted.

From Fisher, hauntography inherits (i) the proactive stance aimed at aesthetic intervention and practical operationalisation, (ii) a *principle of radical alterity* according to which rather than integrating the alien into familiar (and familial) structures, one performs the opposite move, and (iii) the categories of the *weird* and the *eerie*, albeit viewed as affective and fictional modes of the spectral rather than merely of the strange.

The relevance of this chapter is twofold. First, it connects Derrida's and Fisher's hauntological thinking, showing how the *weird* and the *eerie* embody crucial hauntological notions: the Fisherian *weird*, understood as the *out of*, overlaps with the *out of joint* of Derrida's hauntology; the Fisherian *eerie*, understood as a failure of presence or failure of absence, overlaps with Derrida's hauntological reflections on the ghost's vivifying force, the *work* (and thus the agency without presence) that characterises the *Thing* (see Section 6.3.4). Second, the chapter develops hauntography in order to operationalise the hauntological notions discussed throughout the chapter both creatively and analytically: analytical hauntography registers the affective encounter with the *weird* and the *eerie*, and analyses it through the constructs offered by hauntography (see Section 6.2); creative hauntography applies the *weird* and the *eerie* as modes of fiction in order to reveal the spectrality inherent in all things and phenomena. Together, these two hauntographic modes configure a circle of praxis in which practice informs theory and vice versa.

Crucially, as clarified from the outset, hauntography is a framework and not a method (see Section 3.2). This allows one to apply the theoretical tools it offers in diverse fields, from science-technology studies to anthropology, from media and AI studies to cultural studies, from artistic to design practice.

The next chapter applies this framework analytically to music AI through thematic analysis of interviews following a series of collaborative encounters with artists gravitating around the Intelligent Instruments Lab. Through a cross-examination of common hauntographic themes, I aim to respond to my third research question concerning how findings in the intimate human-technology relationship of musical practice aid us in understanding the general impact of AI technologies on our society.

7 AI Hauntography

*Ideas, like ghosts, must be spoken to a little
before they will explain themselves.
Charles Dickens (1847).*

7.1 Introduction

Imagine reviewing a paper for an upcoming conference. Or a student's final essay. Or a PhD dissertation, such as this one. You have been enjoying the writer's voice throughout the text, the stylistic traits that reveal their presence on the other side (of the world, of time, of both). As Derrida puts it, the voice that "marks the name or takes its place" (Derrida 1994a, 9). You get a strange and unsettling feeling as you read on, as though the presence that seemed so distinct in the opening lines had been replaced by another one, much like with the interplay between Helen Harper and the Voder in the dissertation's prologue.

You halt at a sentence that seems slightly off, and go through it a few times. The more you read it, the less it makes sense. Yet, it sounds *as if* it made sense. This leads you to question the line before, and then the paragraph before that, as the suspicion of dealing with a large language model's output grows with every new marker word, such as *tapestry*, *delve*, *intriguing*.⁹⁴ And with it grows your frustration at realising that you might be spending more time in reading the article than its presumed author might have spent in writing it. But this is just an eerie feeling, of a voice within a voice, the trace of something that "works through." You know you will never fully disambiguate this sensation: ghosts, as we saw, are receding and deferring entities, and their spectres point towards the unknown.

We are now in the final substantive chapter of the dissertation, and the theoretical framework to investigate how music AI (and from there AI as a broader phenomenon) exhibits hauntological features has been laid out. In this chapter, I apply hauntography through deductive thematic analysis of interviews with six artists I collaborated with during the past three years. The coding, analysis and discussion of the findings (see Section 3.5), informed by

⁹⁴ LLMs notoriously overuse specific words. These may be used to reveal whether a text is AI-generated or not (see for instance Juzek and Ward 2024).

the insights I gained through the design and artistic practice I described in Chapters 4 and 5 and by countless conversations with colleagues, aim at answering the third question that this research asks: *How do findings in the intimate human-technology relationship of musical practice aid us in understanding the general impact of AI technologies on our society?*

The artists selected for these interviews are a good sample of the diverse group of practitioners I collaborated with as a member of the Intelligent Instruments Lab. In deciding whom to interview, I excluded artists who are aware of and acquainted with my framework, and especially those with whom I engaged theoretically during the collaboration, so as not to risk collecting overly biased data. I also sought to include a diversity of artistic practices, both in terms of interaction design and in terms of type of architecture being applied: while many participants used NAS tools in different configurations, here I include one interview featuring an AI instrument based on Victor Shepardson's Notochord (see Section 2.2.1).

I also tried to collect contributions featuring different modes of engagement with the training data, since findings from Chapter 5, and especially from *End of Empire* (see Section 5.3) and *Mouja+* (see Section 5.4), showed how shifts in this dimension contribute to different understandings of generative models' operations. Finally, I deliberately involved experts in music performance and composition with AI, but also artists with no specific background in this domain, including one who identifies their work within the Historically Informed Performance (HIP) approach (see Section 7.2.4).⁹⁵

Based on a hauntographic analysis of these interviews supported by my own practice, I then discuss three key insights emerging from the interviews, both separately and in relation to each other. First, I look at AI interfaces through the hauntographic metaphor of the *threshold* (see Section 6.5.1). I then reflect on what infiltrates and haunts through these thresholds, and introduce the notion of *spectrophonia*. This concept, paralleling R. Murray Schafer's notion of schizophonia (see Section 2.5.2), describes the idiosyncratic way in which AI pulverises schizophonic sounds, and how these infiltrate the model's outputs like hauntographic ghosts, as per the definitions I offered in Section 6.2. Finally, I continue by investigating another hauntographic manifestation described by the participants, wherein the model is endowed with magical, fantastical and anthropomorphic properties. I thus propose a hauntographic perspective on Marx's notion of commodity fetishism (Marx 1887/1995), and offer an analysis of AI as a *spectral commodity*, which I describe as structurally incapable of concealing the social relations it reifies.

⁹⁵ For a thorough description of the methodology, refer to Section 3.5.

7.2 Hauntography Interviews

As an expansion on the context described in Section 3.5, it is relevant to emphasise that among the tasks of my work package within the Intelligent Instruments Lab’s research project, a crucial one has been that of assisting artists in extending their practices through hardware and software tools, many of which involve the use of AI.

The Organium, the Intelligent Instruments Lab’s library of technical elements for prototyping and thinking through design (Magnusson et al. 2024), provided the raw materials and concepts I used to assist artists and practitioners in developing, among others, feedback systems, electromagnetically actuated instruments, light-controlled flutes, motorised hurdy-gurdies and AI sonic batons (Figure 7.1). As higher abstractions of the Organium’s technical elements, Stacco and Thales found their way into the toolboxes of many artists as well, who used them to produce music for visual media and theatre, to compose for small ensembles and orchestras, and to perform new pieces in live settings. The interviews span applications of my interfaces (e.g. Sections 7.2.1 and 7.2.2), of AI-instruments developed by colleagues at the Intelligent Instruments Lab (see Section 7.2.3), prototypes based on the Organium (Section 7.2.4), software instruments (Section 7.2.5) and interfaces designed by the artists themselves (Section 7.2.6).




Figure 7.1. Instruments developed at the Skerpla composition course. Students from the Iceland University of the Arts (LHI). Clockwise from the top-left: Stretch sensor in Baldur Snær Bachmann's pipe feedback instrument, Karólína Einars Mariudóttir and a light-actuated flute, Daniela Isabel Amado's feedback glass, Konrad Groen's mechanical hurdy-gurdy, and Majella Clarke's augmented baton.

The interviews consisted of individual sessions lasting around one hour each, most of which (except for one discussed in Section 7.2.1) were carried out online. The format was semi-structured, beginning with six baseline questions around the background of the participants, their collaboration with the lab and the role of AI in their practice, followed by a second unstructured part focusing on the themes emerging from the first half of the interview, but also on themes that had emerged through practice before the interview (the questions are listed in Section 3.5.3). I transcribed the interviews and applied deductive thematic analysis as detailed in Section 3.5.3. The resulting insights, sustained by reflections from Chapters 4 and 5, establish the foundation for the discussion that follows.

7.2.1 Hauntogram #1: Þórunna Björnsdóttir's Un-worldings

Þórunna Björnsdóttir is an Icelandic interdisciplinary artist with a background in classical music (piano and oboe) and sonology, currently focusing on electroacoustic music, soundscape composition and contemporary art. Her collaboration with the lab centred on Stacco, which she chose to produce the sound world for *Ashes*, a theatrical piece by Elfriede Jelinek that was performed in Cologne in April 2025. Hauntographically enough, the play juxtaposes the individual dimension of the trauma caused by the loss of a loved person (Jelinek's husband) with the global perspective of the end of the world as we know it (Figure 7.2).

Paralleling the play's structure, which revolves around the four elemental forces (fire, earth, air, and water), Þórunna collected four datasets which I used to train as many RAVE models, including recordings of the Icelandic environment (glaciers, river, and caves) and of Þórunna's voice on extended techniques. Once trained, we mapped each model's decoder input to the readings from a single attractor in Stacco, starting from the second latent (as each attractor in Stacco provides three data points (see Section 4.4.1), these were mapped individually to latents two, three, and four). Þórunna's approach to sound design consisted of timbre transferring her voice through the RAVE encoding function, and simultaneously influencing the resynthesis parameters by biasing the latent dimensions through Stacco (see Section 2.2 for a description of this method).



Figure 7.2. *Ashes*, Elfriede Jelinek (2025). Þórunna Björnsdóttir and I trained a RAVE model on each of the four natural elements, here represented by as many actors.

The interview with Þóranna oscillates between a phenomenological report on her experience with Stacco and reflections on the discourse around AI as a cultural phenomenon, offering interesting insights on how general assumptions around AI manifested locally in her practice. The open-ended part of the interview is particularly instructive, beginning with a statement on how she “took Stacco for what it is,” without thinking “about AI per se” or intelligence tout-court. Nonetheless, she also explicitly acknowledges the instrument’s “intelligence,” which “is there,” “amazing,” but fundamentally opaque and unclear in “what’s behind.” Þóranna attributes this opaqueness to several factors, such as the complexity of the technology, the discourse surrounding data extraction, which negatively biased the actors she worked with (“people are really feeling that they are kind of skinned”), and the overly complex way of explaining the technology which should instead be streamlined through an embodied, “phenomenological” perspective.

As the interview progresses, Þóranna describes the play’s script transitioning from “a real setting” into one folded around trauma, where things “fall apart” due to grief’s undoing. Grief is described as a “world within worlds” rather than an emotion, where time “becomes abrupt” and its flow “is not existing.” She then adds: “For me, Stacco is also like this.” This is due to (i) the rearrangement of the sound material in Stacco’s AI models and (ii) Stacco’s going against the user’s expectations, its playing “a twist upon you,” in that “nothing is ever the same” due to the system’s interaction design and the models’ complexity, which makes the sonic outcomes not entirely reproducible.⁹⁶

The parallel with the Fisherian (and by extension justified in Chapter 6, Derridean) *weird* is evident here, and even anticipated by Fisher himself who describes nested worlds consisting of unresolved spaces awaiting to be reincorporated into familiar symbolic regimes as un-worldings (Fisher 2016, 50). The “awaiting” in Fisher as well as in Þóranna’s parallel between Jelinek’s play and Stacco may be viewed as the measure of a trauma (the severing from one’s reality in Philip K. Dick’s novel discussed in Section 6.1, the ordinary unassimilability of death in Jelinek’s mournings). Stacco is described by Þóranna as an ideal interface to investigate this weird dimension, in that by rearranging the sonic material recorded and selected by the artist in unpredictable ways, it reconfigures one’s world into a strange one. One that resembles the empirical world in many aspects, but also subverts it in more or less subtle ways. As we saw in Section 6.5.1, it is precisely from the intrusion of the alien within

⁹⁶ Once mastered, Stacco allows one to repeat sonic gestures with a certain precision. Here, though, the instrument is used together with timbre transfer from external sources, which highly complicates the interaction.

the familiar that the *weird* emerges. We will return to this in Section 7.3.1 to reflect on how this intrusion plays out broadly in generative AI.⁹⁷

Out of this feeling of a fractured and disjointed time derives a many-layered *eeriness*: first, one inherent to the data used for the training, consisting of vocal growls, long echoes, and sounds evoking natural phenomena; second, in the model’s outputs, where “sounds were so deep and eerie” that the sonic material seemed to inhabit “its reality, its own world somehow” (here the *weird* and the *eerie* appear as completely unified, with the eeriness of the material validating the *weird* – un-beings naturally inhabit un-worlds); third, on the sense of disembodied presence that is produced by the idea that AI is “intelligent and that it’s communicating something back;” Fourth, on “something eerie about our times,” about “the technologies we use and what is going to happen.” These types of eeriness, and the last two in particular, offer an effective operationalisation of the visor effect (see Section 6.2), the gaze we project from what we do not know (Derrida 1994a, 7).

7.2.2 Hauntogram #2: Robert Laidlow and the Out Of

Robert Laidlow is a British composer whose practice has focused specifically on music AI since his PhD at the Royal Northern College of Music’s PRISM centre, in which he investigated AI’s creative and disruptive potential in music composition using a diverse set of symbolic and non-symbolic generative systems. Robert continued his investigation as an Oxford Fellow, and is now AI+ Fellow at King’s College, London.

I collaborated with Robert for a workshop focusing on turning objects of everyday use into AI musical instruments. Robert subsequently joined the lab for an exploratory residency to experiment with Stacco, which led to the inclusion of the instrument in his piece *Techno-Utopia*, composed for the BBC Philharmonic Orchestra with soloist Zubin Kanga.⁹⁸ On this occasion, we trained four RAVE models on the archives of the BBC Orchestra, one featuring the full ensemble, the other three the strings, brass and percussion sections. The system we developed, and to which my colleague Victor Shepardson contributed decisively, features pure latent navigation through Stacco’s attractors, and was performed on the occasion of the May 2025 BBC recording of Robert’s *Techno-Utopia* (Figure 7.3).

⁹⁷ This subversion of the familiar according to machinic principles emerged in other interviews as well (e.g. Section 7.2.2 and 7.2.4) and will be further investigated in 7.3.

⁹⁸ <https://www.zubinkanga.com/>



Figure 7.3. Recording of Robert Laidlow's *Techno-Utopia*. Zubin Kanga on *Stacco* with BBC Orchestra.

As Robert states during the interview, his goal with AI is that of finding “interesting new sounds and new ways of organising sounds,” as a means for a type of artistic engagement where “you shouldn’t need to know that there’s AI involved.” Echoing Þóranna, Robert recognises how these systems are often ascribed a sense of “mystic agency” due to their complexity, something he is interested in exploring creatively, but that he also distances himself from as a composer (“I definitely do not feel that I’m working with a collaborator”).

Throughout the interview, the *weird* and the *eerie* emerge mainly as a byproduct of the friction between the algorithmic and the human ways of thinking about time, space, and memory. Whereas in Þóranna the weird dimension of *Stacco* parallels grief’s un-worldings, in this interview the *weird* (of *Stacco*, but also of how he broadly thinks of AI in his practice) maps more effectively with the *out of* time and space. The *eerie* emerges instead from a sort of listening from the outside-in, which offers “an alien view on what sound is.”

Indeed, in a move that parallels hauntography’s principle of radical alterity (the alien infiltrating the human rather than human categories being projected onto the alien), Robert recognises that AI’s creative capacity stems from the friction of a “completely non-human perspective [...] forced to confront a human perspective”. Whereas music exists in time, “when

you train an AI, it doesn't really sort of know what time is," and "the music is abstracted into a set of atemporal relationships between data points." In other words, algorithms "just are," and the *weird* emerges as a function of our imposing time on them.

By "just being," AI collapses temporal relationships into an abstract (latent) space that one inhabits and explores, provided one has the right tools. Indeed, Stacco is described by Robert as one particularly effective tool, since it offers "a different view of how to organise constituent parts of a piece over time."

As anticipated above, this algorithmic weirdness foregrounds the *eerie*. The architecture's reconfiguration of data beyond time and causality translates into a reorganisation of presence itself. In a reflection that expands from Stacco into AI as a whole, Robert argues that AI accentuates the ongoing spectralising processes of media technologies (as an expansion of this, Robert discusses Facebook slowly becoming a graveyard rather than a social media platform), something that the growing application of AI to interact with lost loved ones has brought to a paroxysm. But, Robert asks, "if you construct a sort of persona using all the data you have of that person, what does that represent or what doesn't it represent? [...] What's lost in there?" This question of digital presence and absence is what prompted Robert to brainstorm with me on a piece in which Stacco is used to access the memories of a digital self (see Section 4.7.1).

7.2.3 Hauntogram #3: Enrique Hurtado and the Cat

Enrique Hurtado is a txalaparta player from the Basque Country, Spain. The txalaparta is a traditional Basque idiophone consisting of a variable number of planks typically beaten by two people using wooden sticks (Figure 7.4). This conversational and percussive nature makes the txalaparta an ideal platform to experiment with architectures that take on the role of one of the two players. Indeed, Enrique's 2015 PhD research focused precisely on designing a conversational system based on Markov Chains to perform with, while his collaboration with the Intelligent Instruments Lab extended this human-algorithm interaction from GOF AI to deep learning using Notochord (see Section 2.2). While Enrique collaborated primarily with my colleague Victor Shepardson, I interviewed him seeking to extend the hauntographic analysis beyond NAS, and into a domain (that of symbolic generative systems) I did not focus upon directly in these three years.



Figure 7.4. Enrique Hurtado in Djúpavík. Performing on the txalaparta with Victor Shepardson's Notochord (2024).

For this project, Notochord was fine-tuned on several hours of txalaparta playing recorded by a group of eight txalaparta players, using contact microphones to capture onsets, velocity, plank ID, and stick being used to hit the plank. In the interview, Enrique describes the collection of the data as a challenging process, in particular due to the participants' wariness in releasing their personal information and "voices" for the training of an AI model. "Everybody is really afraid that the data is going to be used by Google" and people are "rightfully careful." While this is not explicitly described as *eerie*, it nevertheless parallels Þóranna's acknowledgment of unseen cultural and political forces oozing from beneath and beyond the model's hidden layers. This type of eeriness echoes Fisher's own reflection on (harbour) infrastructures and capitalism, in which he argues that the latter thrives on a perception of dematerialisation, while at the same time being materially inscribed, in its continuous shifts of capital and labour, in our technical environment (Fisher 2016, 64).

Enrique becomes more hauntographically explicit as he literally describes the txalaparta-Notochord system as a "weird" and unpredictable actor. However, this weirdness merges here with the sense of eeriness evoked by the disembodiment of the human collaborator into a capricious algorithmic form, the failed presence and absence of an entity that, "like a dog or a cat," unpredictably alternates communication and indifference, participation and detachment. This feeling of "trying to talk to somebody who didn't want to talk at all with you," but then "comes back to communicate, and then again, just go away" is quite different from

the one evoked by Enrique's previous system based on Markovian processes, which felt like "it's making a big effort to try to follow you somehow. Maybe too much, even." In other words, where GOF AI seems to miss the sweet spot of an engaged interaction due to a complete absence of volition, AI is here perceived as if exhibiting too much volition, as if pursuing a different and unknown agenda, choosing when to care and when not to care, when to establish a channel of communication and when to cut it.

While we might regard Enrique's metaphor of the model acting as a cat or a dog as an anthropomorphising statement, hauntography's principle of radical alterity invites us to look at how the biology Enrique refers to differs from ours. The communication Enrique describes is indeed an interspecies one, and as such it inevitably entails an extremely high degree of misunderstanding. Enrique is unsure whether the txalaparta's temperamental personality might be determined by the AI model or whether it results from other variables (for instance, changes in the gain levels in the piezoelectric sensors, or in the tracking algorithm), but the point, as we saw, lies precisely in not-knowing. Under this light, the phrase "a cat or a dog" represents Enrique's tentative answer to Walshe's question of "who is listening, and how?" (Walshe 2023).

Using Fisher's discussion around Nigel Kneale's *Quatermass* (Fisher 2016, 83), we can push hauntography's principle of radical alterity even further, and with it the eeriness it foregrounds. Fisher describes how Kneale's 1952 *Quatermass* series "discovers the mechanical-automatic-alien basis of what has been taken to be human," due to how the distinction between the player and what is being played breaks down (Fisher 2016, 87). The *erie* emerges out of the realisation of "external" forces within, which implies that humans might be less present to themselves than they are willing to admit, inextricably caught in patterns of mechanical repetition.

A similar reversal emerges in Enrique's interview as well. The predictable and repetitive behaviour of his previous Markovian system allowed him to infer a clear boundary between the roles of the human and of the machine, with the former proposing and the latter following in a relatively predictable way. With the new Notochord-based system, the one who follows is instead Enrique, who has "to play very, very stable." Enrique strongly stresses this aspect, saying: "if I try to do something very very weird, it would be [gestures emphatically]." This inverts the paradigm, with the human caught in "doing the same thing," a repetitive, almost machinic pattern, as the AI system's zoomorphic apparition "just [goes] somewhere." The eerie conjecture that haunts Fisher, of life being a mere region of death, is foregrounded here by the twofold question of who or what is playing, and of who or what is being played.

7.2.4 Hauntogram #4: Darija Andzakovic, the Brainchild and the Memory Palace

Darija Andzakovic is a double bass player specialised in historically informed praxis, and a musicologist from former Yugoslavia, currently pursuing a doctorate in historical musicology focused on the eighteenth-century double bass virtuoso and composer Johannes Matthias Sperger. I collaborated with Darija's BYOV (Bring Your Own Viola) duo, which includes my wife Natalia Duarte Jeremias, on two occasions, both focusing on the second movement of Sperger's 1789 sonata for viola and contrabass, which he wrote while travelling from Trieste to Venice.⁹⁹ During our first collaboration, which predates this PhD research, we retraced Sperger's route through land and collected field recordings and images from the areas. I then combined the material into a fixed-media study including the duo's performance of the sonata.¹⁰⁰

Once my collaboration with the Intelligent Instruments Lab began, we decided to further develop the sound material we had collected into a performance premiered at the 2023 AIMC Conference (Andelic-Andzakovic et al. 2023). For the occasion, I trained a RAVE model with a curated selection of field recordings from the database we had collected. I then designed two wireless wristbands combining a gyroscope and an accelerometer, one for Darija and one for Natalia. The wristbands, mounted on the players' bowing hands, were mapped to RAVE latents in the decoding phase, from latent seven to latent twelve for a total of twelve. I instead used Thales to control the first six latents live. The performance consisted of an improvisation over the structure of the second movement, where the sounds of the Italian northern coastline and melodic fragments from Sperger's sonata merged (Figure 7.5).

99 "Trieste" appears at the start of the manuscript, indicating that the sonata was being written there. "In Venezia" appears at the end of the manuscript, indicating that the parts were copied or composed on the way to Venice (or perhaps completed there).

100 https://youtu.be/sUE_hlfU074?si=nPqG04jzMXk7Au_



Figure 7.5. BYOV performing over Sperger's sonata at AIMC 2023. The red line navigates the score based on the performers' gestures, suggesting different melodic fragments.

Darija's interview proved dense with hauntographic cues. While the first collaboration is described as a sort of "literature review" in a thesis writing, the performance piece left Darija with the impression of coherence of a finished thesis, where "we created something out of everything that we'd gathered." In this latter phase, what struck Darija was how we outsourced memories precisely in order to "continue to experience them and react to them." Darija adds that "once all of those sounds were there [...] we got new ideas." Here, we encounter again an operationalisation of the principle of radical alterity, where AI does not follow a logic of anthropomorphisation (as Darija puts it "say, a virtual girlfriend that is slowly destroying your ability to interact with the world"), and instead is "somehow adding to that memory palace." This notion of a "memory palace" identifies neither the architecture nor the dataset. Rather, Darija uses this term to identify a conceptual construct, a "latent space," one might say, of subjective memories she had developed in the experience of collecting the field recordings, and of which the model's algorithmic rearrangements "kind of changed the curtains."

As soon as memory is outsourced and reconfigured, an eerie space opens up. The "memory palace" turns into what Darija describes as a "brainchild," an assemblage in which "all the different kinds of artistic genetic materials were put together" into "something new that goes on its own." As the interview progresses, Darija discusses in a mockingly mournful tone

how the first “baby” we trained (a model we discarded due to an overly noisy output) had to be “locked in the basement,” since it was giving “Hellraiser” vibes. “It felt like someone had opened the puzzle box” from the movie, unleashing the demons.

As the interview approaches a conclusion, Darija weaves together many of the spectral threads scattered all around. What seems to tie coherently the new curtains of the memory palace and the eerie vibe of the brainchild is the subjective experience with the data used for the training. The places we (and Sperger) visited and the visiting itself established the necessary frame of reference, the scaffolding and organising principle of the model’s meanderings, thanks to which memory, even though outsourced, pulverised, and reconfigured into something eerily other, could still be understood as such (the furniture in the memory palace, the memory genes of the brainchild). This strongly resonates with my experience in working on *End of Empire*, and with how my understanding of the model’s output changed as I got acquainted with the experiences and stories that led to the collection of the data used to train it. This insight will be central to the discussion in Section 7.4.

7.2.5 Hauntogram #5: Erik Klinga and Sonic Hauntology

Erik Klinga is a Swedish musician and composer, known for his activity with the Swedish indie band Simian Ghosts. Erik’s formal education is in electroacoustic composition and music performance. We met during his residency in Reykjavík, where he worked live with a Buchla system. During our collaboration, I fine-tuned a RAVE organ model from the Intelligent Instruments Lab with a dataset Erik collected out of three of his albums, including organ textures and synthesizer drones. To integrate the model into Erik’s setup, I designed a Max for Live device that allows for the individual manipulation of the decoder’s latent dimensions or for timbre transfer through the encoding function, or for a combination of both, with a signal (in Erik’s practice, typically from his modular synth) being timbre transferred while Erik offsets the latent states with a MIDI controller.

After experimenting with the device and model for a few months, Erik’s feedback has been largely positive. According to Erik, the NAS architecture’s autonomous organisation of entangled sounds into the latent space (see Section 4.2.2) does not afford a “straightforward linearity,” to the point that the model needs to be interacted with and “investigated sort of on its terms.” This unpredictability is very welcome to Erik, who enjoys being surprised and steered towards unexpected musical discoveries by the technologies he uses.

At our first meeting, in witnessing Stacco's vocal grammelot for the first time, Erik mentioned Fisher's sonic hauntology, a domain of artistic investigation he said he had written extensively about (unsurprisingly, given the name of the band he toured with for more than ten years). During the unstructured part of the interview, I asked Erik to elaborate on this. His response initially focuses on the uncanniness of the vocal reconstructions, impossibly human and thus eerie ("when it is AI and its voice and it's in between, it becomes almost like that uncanny valley situation"), but then continues discussing the "distance effect," a phrase he uses to identify the dreaminess introduced by multiple and successive layers of mediation, introducing glitches, distortion, and traces that, independently of any intelligent behaviour, turn the model into an agent that actively shapes the music.

Yet, fragments are extracted, disassembled and reassembled based on a logic that responds to new rules, and that surprises and inspires. According to Erik, these processes are not merely nostalgic in Jameson's sense (see Section 6.4), but rather introduce something entirely new: the model, argues Erik, augments the material "in a sense that it sort of becomes something new extracted from it," it "continues into the next, in a way."

The notion of nostalgia, as we saw, identifies the postmodern folding of culture onto itself, a stereotyped form of montage that displaces the subject outside of time (see Section 6.4). Montage, according to Fisher, also represents a quintessential form of the *weird* (Fisher 2016, 11). But there exists a crucial difference between nostalgia and weird montage: while the former is apprehended as caricature, the latter, which sonic hauntologists developed and applied as a reaction to the former, reorganises the past in order to recover the very possibility of a future, or as Erik puts it, of "continuing into the next" (see Section 6.4.2 for a discussion of the difference between the two).

7.2.6 Hauntogram #6: Voldemars Johansons and the In-between

Voldemars Johansons is a Latvian electroacoustic composer and researcher teaching at the Latvian Academy of Music, whose interests include algorithmic composition and sound spatialisation. I met Voldemars on the occasion of the SWRL Festival he organises annually in Riga, where I premiered *Mouja+* and ran a workshop on NAS in music composition (see Section 5.4). During the workshop, I taught a group of composers how to train RAVE, the best practices for data curation and a series of strategies to integrate NAS in one's workflow. Voldemars trained a model on environmental recordings from the Arctic and Nordic regions,

including ice formation, geothermal activity, and other natural phenomena he documented during a research trip to the Svalbard islands. He then used the model to compose an acousmatic fixed-media piece he presented at SWRL, as part of a cycle titled *Songs of Disappearance*, addressing climate change in boreal environments (Figure 7.6).



Figure 7.6. Voldemars Johansons at SWRL. Playback of Johansons' piece from the 2024 cycle *Songs of Disappearance*.

In the first part of the interview, Voldemars focuses on the technical challenges he encountered in working with RAVE. Similarly to Erik, he became highly aware of the “non-linearity” of the model, which introduces a welcome element of unpredictability, which in the case of Voldemars also entails some frustration in “finding the sweet spot” within a latent space often populated by noisy artefacts.

When Voldemars describes NAS as “the least straightforward tool for synthesis” he has used so far, this is not necessarily to criticise it, but rather to emphasise the novelty of the technology, and to point to the fact that he “couldn't compose a similar piece with any other means.” This translates into what Voldemars describes as the “weird control” of the system, a notion that describes the model’s refusal to respond like any other previous synthesis method, in how it requires reframing one’s understanding of and approach to algorithmic music. This weirdness is about shifting one’s focus from screens, keyboards, and programming languages,

and embracing a process of resistance and surprise that is in many respects “similar to working with a physical material.”¹⁰¹

In the interview, Voldemars discusses NAS mainly from a technical angle. Yet, toward the end of the interview he also reflects on how RAVE’s affordances algorithmically translate his own aesthetic stance, offering a hauntographically relevant insight. Here, Voldemars discusses the “in-between spaces” he aims to capture sonically: for instance, the transitioning of natural phenomena from one state into the other; or their disappearing and appearing as something new, which his cycle *Songs of Disappearance* tries to capture. As he argues, the synthesis approach afforded by NAS, where data transitions into something “very different” that nonetheless retains “echoes of the source material,” a “reflection feedback of the source,” represents the ideal method to musically depict the exponential changes of the Anthropocene, wherein the ecosystem bears the traces of eerie agents too large to embrace and the remnants of natural processes that are no more.

7.2.7 Emerging Themes

The hauntographic reading of these interviews, along with local ones, revealed a series of overarching themes that will be developed in the following sections. First, AI systems are often described as liminal gates (as I describe them in Section 7.3, *weird thresholds*) that afford and complicate the relationship with the data. Indeed, for some of the participants, the models estrange memory, making it walk on algorithmic legs while bearing the echoes of the source material (Darija’s term *brainchild* aptly exemplifies this). Second, there is a certain resistance of models and interfaces to respond as expected, sometimes ghosting the users, sometimes playing “a twist” upon them or “going on their own,” sometimes resisting through algorithmic affordances that translate into material ones, and often requiring to be approached on their terms. Third, some of the interviews report on a certain anxiety and mistrust, with the broad discourse on AI as an extractive tool infiltrating and haunting its local uses. This anxiety motivates the chapter’s concluding hauntographic expansion into AI as commodity form.

101 This parallels findings on how NAS is often described in terms that are more similar to how a traditional musical instrument is approached, as discussed in Section 5.5.3 (Privato, Lepri, et al. 2024) (Zappi and Tatar 2025).

By further developing these points through a hauntographic lens, the following sections address the third research question relative to *how findings in the intimate human-technology relationship of musical practice might aid us in understanding the general impact of AI technologies on our society*. The discussion is divided into three parts. The first section expands on the first emerging theme, discussing hauntographically the notion of *threshold* as applied to NIMEs. From here I expand on algorithmic presence, memory and knowledge transmission in music AI. Finally, in the third section I extend the hauntographic reflection into AI as commodity form, a perspective necessitated by the anxieties around AI emerging from the interviews, by the research question in itself, but also by the political angle that characterises hauntology from its inception.¹⁰²

7.3 The NIME at the End of the World

In *weird fiction*, we are often confronted with thresholds between worlds, argues Fisher (Fisher 2016, 28). Indeed, the *weird* entails an eruption into empirical reality of something from the outside, which subverts the way what is taken as given is interpreted (see Section 6.5.1). And for the outside to infiltrate and turn upside down our world, some kind of passage needs to open, whether material or symbolic, explicitly framed or hinted at.¹⁰³

In *The Weird and the Eerie* we encounter many such thresholds, doors and portals: H.G. Wells' novel *Door in the Wall*, in which a door appears intermittently throughout the life of the protagonist, revealing a dimension he hesitates to set foot inside but nevertheless haunts him in his daily life (Fisher 2016, 26); the time-travel apparatus in Tim Powers' *The Anubis Gates* (Fisher 2016, 40); the signature red curtain in Lynch's *Mulholland Drive*, or the labyrinths in its *Inland Empire* (Fisher 2016, 28). While these gates of the *weird* offer the possibility of an incursion into the outside, they also require leaving one's terrestriality behind (Lovecraft in Fisher 2016, 20). In other words, to step into the unknown, one needs to first welcome it within. Thresholds are thus *weird* insofar as they afford transitions that are never unidirectional, but rather provoke the incursion of external elements, that reveal deeper

102 In *Specters of Marx*, Derrida parallels the "automatic autonomy" of commodities under capitalism with that of "technical life" (Derrida 1994a, 192). Confronting the notion of commodity fetishism as applied to AI appears thus inevitable.

103 For a hauntographic definition of threshold, see Section 6.2.

structures that are always already at work, and of whose existence we can be unaware or oblivious.¹⁰⁴

Throughout the interviews, accounts of Stacco, Thales, Notochord, and other hardware and software systems of interfaces and generative models often entailed a view of the system as a *threshold* device. This angle also emerged in designing Stacco (a Ouija board is indeed a threshold device, see Section 4.7.1), and became the foundation to one of my performances (see Section 5.2). In the interviews, *our* side of the threshold typically implies a dimension of entropy and time, of causes and effects, of matter and thought, of presence and absence: an empirical reality understood as linear. On the other side lies an alien world of memories and traces, an oneiric latent space and “latent time” of correlation without causation, of inaccessible computational processes and “worlds within a world” as in Section 7.2.1, that introduce, as Robert puts it, an “alien view on what sound is.”

The gate separates and connects. Like the mushroom imagined by Anna Lowenhaupt Tsing, which catapults us into the capitalist aftermath only to reveal it at work here and now (Tsing 2021), these AI NIMes stand at the end of the world, functioning as gates for a weird outside that is creatively meaningful insofar as it subverts and reconfigures how the inside is understood, revealing it as always already artificially plain, linear, causal (for instance, think of Jelinek’s world within world of grief as described by Póranna).¹⁰⁵

Yet, while thresholds enable weird incursions, they also represent an informational bottleneck that conceals and restrains communication. Referring to Louise Meintjes, Sterne discusses how the technical apparatus of the studio “procures for those who have earned access to it the services of that force, or spirit, lodged within it” (Sterne 2019). Similarly, the unknown beyond the NIME-threshold reifies as the mystery of the threshold itself. Póranna frames this transfer of powers into the interface quite explicitly when she discusses grief as a “world within worlds” where “the flow of time is not existing,” and adds that “Stacco is also like this” in that “it plays a twist on you.” Similarly, in Robert Laidlow’s practice, Stacco’s “messiness” allows one to dive into one’s memories as reorganised by the model, as if “swimming in them.”

Here, the *weird* foregrounds the *eerie*. In Enrique’s txalaparta practice, the weirdness of the model’s erratic interactions translates into the eeriness of a highly unpredictable agent, which feels all the more present as it withdraws itself from the user. As Sterne argues, “when

104 Fisher often frames these forces through Lacan’s *weird* psychoanalysis.

105 One way of framing this is through what Cristianini describes as one of three AI shortcuts (2021), namely AI’s organisation of the data according to correlation patterns rather than causality.

people do not have access to the inner world of an instrument [...] they are more likely to attribute to it a vast, complex inner world” (Sterne 2019). What inhabits this complex hidden world is anyone’s projection, but not-knowing acts as a vacuum, a void that fills with eerie presences, capricious spectres by which we are seen but that we cannot see nor fully understand. In Enrique’s interview these forces are all the more present as they act unpredictably and volitionally, like “cats and dogs,” becoming all the more present as they disengage the user, as they turn themselves off as a twenty-first-century *Useless Machine* (see Section 6.6.3).

Whether the NIME’s unpredictability results from a complex interaction of design choices, from material affordances or from architectural idiosyncrasies, knowing that some kind of process takes place on the other side, of which we are not fully aware, is enough to conjure the ghost. And one needs to localise it (“One has to know. One has to know it. One has to have knowledge” argues Derrida) if not tangibly, at least through the NIME that marks the name or takes its place (Derrida 1994a, 9). Localising the ghost to conjure it away, and conjuring away the ghost to make it present, herein lies the enjoyment with these threshold instruments. In this act of *exorcism* (see Section 6.2), the absent transforms itself and the world all around; it demands to be acknowledged and even served, and the *weird* world of the NIME-threshold becomes *eerie*.

7.3.1 *Weird Thresholds and Ground Truth*

According to Fisher, thresholds work in different ways in *fantasy* and *weird* fiction: where the former presents us with parallel timelines that never touch or affect one another (here thresholds do exist, but as a sort of wormhole into a completely alien outside), the latter builds its vertigo precisely on their collisions. Weird fiction thus needs to invest heavily in anchoring the “here” and the “there” in some sort of “ground truth:” Philip K. Dick’s meticulous construction of Gumm’s fictional world as a sort of uncanny 1950s theme park in *Time is Out of Joint* (Dick 1984), for instance, serves the purpose of amplifying the weird effect that derives from the infiltrating cues of its fictitiousness.

In analysing the interviews from the perspective of this dichotomy, it is clear that the data upon which the architectures are trained is what is shared between the “here” of our world and the “there” of the generative model, what embodies the “ground truth” that the *weird* necessitates. This is an idiosyncratic trait of machine learning systems, which in music

distinguishes them from other synthesis techniques in which sound is generated by analogue or digital circuitry, or, when dependent on existing material, is rearranged according to a known logic.

I argue that while music always implies constructing and accessing new sonic worlds as does fantasy fiction, AI NIMEs embody the specific mechanisms of weird fiction, functioning as *weird thresholds* that subvert the sonic material, and the world it points to, in algorithmic terms. Stacco and Thales are evidently thresholds to the *weird* for Robert, Póranna and Darija, who use them to engage with an “alien view on what sound is.” Similarly, RAVE became a threshold to the *weird* for Philippe Esling, who in Section 5.2 used it to encounter choirs with no lungs, and in my practice *End of Empire*’s central column welcomed visitors into a weird dimension of unintelligible yet eerily familiar vocalisations (see Section 5.3.4).

In the above examples, the weird subversion discussed above involves both the sonic world represented by the data and the things in the world that sound represents. Thus, establishing a ground truth requires one to establish some situated knowledge of the sources of a certain dataset. For Darija Andzakovic, for example, the “memory palace” was there from the outset (the outset being the sound walk and recording sessions from Trieste to Venice), which the NAS model reorganised by “changing the curtains.” Similarly, Robert Laidlow approached Stacco as a threshold to understand and investigate the BBC Orchestra’s archives. Erik Klinga and Voldemars Johansons also recognised how the model’s output maintains echoes of and at the same time reorganises into new configurations (thus “continuing them into the next,” as Erik argues) the sounds they had collected first-hand. Similarly, my installation *End of Empire* became weird once I situated the data used for the training of the NAS models in my personal experience (see Section 5.3).

In order for the *weird* to emerge, in addition to the subjective ground truth discussed above, the NIME-threshold requires a diverging dynamic, which causes what is well-established and known to become alien. In the next section, I describe this mechanism hauntographically through a parallel with other spectralising dynamics of sound mediation and transmission I introduced in Section 2.5.

7.4 Spectrophonia

In Chapter 2, I described R. Murray Schafer’s notion of schizophonia as the severing of a sound from its source, which thus becomes schizoid, moving in all directions and independently of its origin (Schafer 1969). This notion returned in discussing the *weird* in *Mulholland Drive*’s Club Silencio, where what Schafer describes as the ventriloquising of modern life is progressively magnified to the paroxysmal conclusion of Rebekah Del Rio fainting on the stage with her voice lingering through the room, having lost its newly found embodiment and apparently searching for a new one (see Section 6.5.1).

This effect sets the stage for a similar yet distinct movement that emerged through this hauntographic investigation and from practice. Here, I observe the weird disjuncture of schizophonic sounds, as these are further pulverised, reconfigured and distanced from themselves, and the eerie familiarity in the algorithmic outputs that display the spectral traces of the multitudes they derive from (ghosts in the most hauntographic sense, since they can only be hinted at through the spectres they conjure, and never precisely located). I use the word *spectrophonia* to describe this mechanism of eerie and weird remembering.

7.4.1 Defining Spectrophonia

To better frame spectrophonia, it is useful to compare it with Schafer’s notion of schizophonia, but also to expand it comparatively through Steven Feld’s notion of schizophonic mimesis (see Section 2.5.2). Spectrophonia differs from schizophonia in at least three ways. First, whereas schizophonia focuses on the severing of sound and source, spectrophonia assumes the schizophonic act as prior to the spectrophonic one (a dataset of disembodied sounds is necessary to train a model), and thus should be viewed as a further abstraction from it. Second, whereas schizophonic sounds emerge typically from one localised source, spectrophonic sounds converge from the statistical rearrangements of a multitude.¹⁰⁶ Third, spectrophonia frames AI-generated sounds hauntographically and understands the sources as hauntographic ghosts: vivifying forces that exist in a state of difference and deferral (thus, of *différance*) and

¹⁰⁶ By “convergence” I refer here to the phenomenological experience of hearing traces of multitudes from the dataset collapse into a single sound (see for instance the example in 5.4.7). I do not intend this term as a technical description of the model’s workings, which rather functions by statistical abstraction and dimensionality reduction.

are experienced through their spectres in the model's outputs. Moreover, it should be noted that while Schafer develops schizophonia through a critique of the degradation in the "natural" soundscape, spectrophonia provides a framework for describing AI-mediated sonic processes which may be used both interpretatively and as a tool for an ethical engagement with AI-mediated processes (as further discussed in Section 7.4.2).

During the interviews, participants expressed several ideas that touch on this spectrophonic nature of NAS synthesis. This emerges, for instance, in the anti-nostalgic and hauntological tendency of RAVE's sound to function as montages that nonetheless "continue into the next" as described by Erik Klinga, and in the echoes of algorithmically rearranged sounds in Voldemars Johansons' *Songs of Disappearance*. Similarly, in *Mouja+*, the investigation on sounds of trains and public transportation evoked by the prompt "exit" arguably exemplifies an artistic engagement with spectrophonic dynamics (see Section 5.4.7).

In music AI, examples of spectrophonic processes are ubiquitous. While in the future I would like to fully engage with this notion beyond the case studies of Section 7.2 and my own practice, to further expand the reflection it is nonetheless worth advancing at least one example, in the ghostly appearance of the band Velvet Sundown.¹⁰⁷ This AI-generated rock band rose to notoriety in summer 2025 with a repertoire and aesthetic reminiscent of 1960s bands such as Creedence Clearwater Revival.

At the early stages of the fictitious band's trajectory, listeners were lured into following it through the Spotify playlists and recommendation system. In order to enter such playlists, the band's chord progressions, sound, and production style had to be endowed with enough residue from 1960s aesthetics to register as such. Yet, as soon as we listen more attentively, the music becomes clearly derivative, with vague, meandering lyrics and manneristic formal structures. In the very act of appropriating by algorithmic means the style of a particular group of artists from a particular point in time, the music got haunted by their ghosts.

This dynamic inevitably influences the cultural and social proliferation of sound, which Feld describes through the notion of schizophonic mimesis (see Section 2.5.2). Indeed, due to the algorithmic pulverisation of the data and convergence from a multitude that characterises spectrophonic processes, it becomes increasingly complex if not impossible to trace the different stages of proliferation of sound, and to identify the different negotiations and unequal power structures in which these are performed.¹⁰⁸ In hauntographic terms, with NAS when

¹⁰⁷ <https://www.velvetsundownmusic.com/>

¹⁰⁸ To discuss how sounds circulate through different chains of appropriation, constantly negotiating claims to authenticity, ownership, and cultural meaning within relations of unequal power, Feld offers the example of the

attempting to trace a sound back to its source, one is confronted with the deferring principle of a ghost. Moreover, the AI architecture in itself becomes yet another actor in this negotiation, one that contests the aura of authenticity of the sources it reconfigures. Indeed, now that the hoax has been exposed and the band gained momentum, *Velvet Sundown*'s website frames the band's algorithmic origins as a point of pride.¹⁰⁹

This dynamic of authorial and algorithmic lineage and appropriation will be further investigated in Section 7.5 by expanding the hauntographic reflection into AI as commodity form.

7.4.2 *Spectrophonia in Artistic Practice*

As my artistic practice and the hauntographic interviews suggest, this spectrophonic tension of disjunction and lineage discussed above can be experienced differently depending on the listener's understanding and relationship to the training data. In *Mouja* (see Section 5.2), the hauntographic spectres I engaged with were those of glitches and digital distortions from sparse areas of the latent space, indexical more to the architecture's presence than to the data. In *End of Empire*, by contrast, my personal involvement with the people and places that contributed to the training data afforded an engagement with a different kind of spectres, this time indexical to the people and social relations embedded within the dataset.

Darija Andzakovic could similarly perceive the model's different arrangements of the same "memory palace" precisely because she had walked in person the route from Trieste to Venice, collected the field recordings, and used this experience to understand the historical context of Sperger's journey. The RAVE model's reorganisation of the sounds registered as both familiar and "other" (as a matter of fact, a "brainchild") because she had developed an interpretive frame to recognise the traces in the outputs.

These observations point toward another, more artistically engaged, reason for paralleling spectrophonia and schizophonia, that I identify in the body of works developed by the members of the World Soundscape Project (WSP) for a situated ethics and practice within

Pygmy *Hindewhu* song. Here, Feld traces the chain of mediation of this song from its origin to Herbie Hancock and Madonna, arguing that the song is progressively detached from its social and ecological context (Feld 1996).
109 The Velvet Sundown website now reports: "We are not real in the traditional sense — but our melodies are. Our vibes? Immaculate. Our royalties? Minimal. We exist to confuse music journalists, comfort robots, and help Spotify executives sleep at night. Are we replacing human artists? No. (Yes.)"

the schizophrenic soundscape. Barry Truax summarises the principles of soundscape composition into four points, that I synthesise as follows: (i) maintaining the recognisability of the source material, (ii) invoking the listener's knowledge of the environmental and psychological context, (iii) allowing for the composer's knowledge of the context to influence the shape of the composition, (iv) using the process to enhance the listeners' everyday perceptual habits (Truax 2008).

While these four principles represent a generalisation of the diverse practices of the WSP, they nevertheless offer an overview of how these composers thought of the schizophrenic soundscape in relation to the situated knowledge of the listeners, one that might arguably be adapted to the spectrophonic soundscape that music AI foregrounds. For the WSP, engaging with the new soundscape implied fostering the listener's awareness of the sources and the processes of technical mediation. In the case of a spectrophonia-informed practice, the principles of soundscape composition may be translated into a practice of cultivating a certain awareness of the social relations, of the bodies and events encoded within and spectrophonically deferred by the training data.

7.5 (Another) Spectre of Marx

The third and final reflection offered in this chapter aims to expand the discussion into AI as a spectral phenomenon under capitalism. I consider this theoretical expansion necessary for a series of reasons, not least the preoccupation of the artists I interviewed, who often criticised AI as an extractive instrument, expressing a concern that regularly emerged in conversations on the topic throughout the research, and that I confronted also in my practice (see Section 5.3). As Enrique effectively put it, “everybody is really afraid that the data is going to be used by Google.”

This preoccupation manifested in various ways throughout the interviews. This is evident, for instance, in Póranna's stance toward AI, seen at the same time as a projection of corporate interests (“how you feel you're being robbed”) and a creative agent of *weird* spectralising powers, that plays a “twist on you.” Enrique also had to confront a healthy dose of mistrust, as he struggled to convince the participants in his study that their contributions to the dataset would serve research purposes alone. Darija's interview is particularly indicative here. By recognising AI's potentially exploitative nature, describing how some forms of AI

“slowly suck your soul,” while at the same time acknowledging its creative possibilities in performance practice, Darija exemplifies an ambivalence that, beyond these interviews, permeates the discourse around AI.

The following discussion also serves a methodological and a structural purpose. From a methodological standpoint, the aim is to demonstrate the potential of hauntography as an instrument for cultural analysis. From a structural standpoint, the aim is to provide a circular closure to hauntography’s political trajectory, which is integral to Derrida’s and Fisher’s hauntological critique, and that hauntography borrows from these thinkers. Scholarly literature after *Specters of Marx* has focused extensively on spectrality, and engaged perhaps less in depth with the hauntings of the Marxist spectre. By conjuring the latter, this section offers a materialist perspective on hauntological thinking in which the spectrophonic hauntings discussed in Section 7.4 become evidence of much deeper, structural ones, relative to the AI commodity.¹¹⁰

7.5.1 *Commodities and Spectres*

As seen in Chapter 6, Derrida’s and Fisher’s hauntological reflections rest, among other things, on a view of capitalism (and neoliberalism) as a spectral phenomenon. The views of these thinkers might be briefly framed through a hauntographic perspective as well, that is, by looking at capitalism as inherently *weird* and *eerie*. If we apply these categories to Derrida, capitalism appears *weird* in its covert attempt to install itself out of time, as natural law: the temporal disjuncture wherein the death of the Marxist spectre is announced precisely in order to exorcise its return. It appears instead as *eerie* (at least) in the performativity of this exorcism, which betrays the covert agency of the political (even holy) alliance that performs it. For Fisher, capitalism is *weird* because it places the subject, rather than itself, out of time, stuck in the disjunctures of a future that projects into a nostalgic past. Finally, Fisher explicitly describes capitalism as *eerie* in that it attempts to appear as a force that does not exist “in any substantial sense, yet it is capable of producing practically any kind of effect” (Fisher 2016, 28).¹¹¹

110 In *Specters of Marx*, Derrida writes: “Marx does not like ghosts any more than his adversaries do. He does not want to believe in them. But he thinks of nothing else” (Derrida 1994a, 57). While Derrida’s critique of Marx would deserve a hauntographic analysis on its own, it is relevant to stress that spectrality represents for Marx a crucial indicator, the locus of the commodity, where one encounters the upside-down world of capital.

111 The words “attempts to” are relevant here. For Fisher, capitalism “wants” to be understood as a hyperobject, but it is, in fact, a quite mundane one, whose features are locally inscribed in all of the global infrastructure.

Here, I would like to add to the two hauntographic perspectives above another perspective informed by one of the Marxist spectres (indeed, as Derrida admits from the very title of the monograph, there exist many of them). Rather than from the opening lines of the *Communist Manifesto*, the one I conjure here haunts the reader from the first chapter of *Das Kapital*, and deals with the fetishistic nature of commodities and the “secret” thereof. Here, Marx writes:

It is as clear as noon-day, that man, by his industry, changes the forms of the materials furnished by Nature, in such a way as to make them useful to him. The form of wood, for instance, is altered, by making a table out of it. Yet, for all that, the table continues to be that common, every-day thing, wood. But, so soon as it steps forth as a commodity, it is changed into something transcendent. It not only stands with its feet on the ground, but, in relation to all other commodities, it stands on its head, and evolves out of its wooden brain grotesque ideas, far more wonderful than “table-turning” ever was (Marx 1887/1995, 47).

One could not hope for more hauntographic markers in a single paragraph. The “transcendent” nature of the commodity foregrounds the transformation of the products of labour under capitalism into *weird* objects, “grotesque” artefacts (notably, the grotesque is yet another manifestation of the *weird*, according to Fisher (Fisher 2016, 32)), that stand on their head (literally *out of place*) in relation to all other commodities. As Marx argues in a note appearing in the German edition, the weirdness of this transformation parallels (and even surpasses) that of “table-turning,” the séance practice that became particularly fashionable among the European aristocracy of the nineteenth century (Marx 1887/1995, 56).

The weirdness of the commodity that stands upside-down foregrounds here an element of eeriness: the table-commodity has a “head,” and even a “brain,” and thus a volition, exemplifying how in capitalism “the productions of the human brain appear as independent beings endowed with life, and entering into relation both with one another and the human race” (Marx 1887/1995, 48). This dynamic points to the particular form of fetishism that characterises the commodity under capitalism, whereby commodities obscure and reify their origins due to the unplanned and thus opaque dynamics of production, and are thus prone to assuming mystical character and powers. It follows that “the whole mystery of commodities, all the magic and necromancy that surrounds the products of labour as long as they take the form of commodities, vanishes therefore, so soon as we come to other forms of production” (Marx 1887/1995, 57).

We do not know of the commodity's origin (not because it becomes opaque through use as in Latour, but because of the particular conditions of capitalist production), and thus we endow it with (or it seems to assume in itself) certain properties based on how it dynamically enters into relation with the human race. Of these properties, exchange value is a crucial one, but not the only one. As Derrida notes in a passage referencing Marx's example, the table stands and turns itself upside down (the emphasis here is on its ability to act). It is ambiguously autonomous and automaton (Derrida 1994a, 227): a strange creature, at the same time "Life, Thing, Beast, Object, Commodity, Automaton—in a word, specter" (Derrida 1994a, 190). Enrique Hurtado's zoomorphic AI-creature, Darija's "brainchild," and even the subtle gaze from the outside described by Robert and the AI trickster described by Þóranna, thus subsume within the vivifying principle of the spectre the paradoxical agency of the commodity as fetishised labour.

7.5.2 *AI as a Spectral Commodity*

In generalising hauntographically from the case studies proposed in this chapter to AI broadly intended,¹¹² it is crucial to begin by distinguishing the fetishistic spectrality of the commodity form from the spectrophonic hauntings discussed in Section 7.4. Indeed, the latter describes how given certain conditions the social relations that constitute the data used to train the model infiltrate its outputs. The former, on the other hand, represents a reification of the model-agent that individuates it *outside of* these very social relations.

The haughtograms discussed in Section 7.2 exhibit a spectrum of positions around the two haughtographic angles described above. While Robert Laidlow dismisses the idea of the model-as-collaborator to focus on how it offers an outside view on archival memory, Enrique Hurtado seems less interested in the spectres conjured by the data, and focuses on the zoomorphic appearances and behaviour of the model. Voldemars Johansons and Erik Klinga offer a slightly less antithetical perspective on this, in that they both endow the model with a certain agency (although one that does not entail presence), and at the same time reflect on how, through it, they could experience their data anew. For Þóranna Björnsdóttir, the eeriness of the generative system comes from a certain perception (and assumption) of intelligence, and

¹¹² While this section broadens the discussion from the specific architectures and systems developed at the Intelligent Instruments Lab, the acronym "AI" continues to refer specifically to the recent wave of deep learning systems rather than GOF AI ones, as detailed in Section 2.2.1.

at the same time from the data's own eeriness. Similarly, through the revealing metaphors of the memory palace and of the brainchild (here, it is impossible not to think of a certain table with a certain head and a certain brain), Darija Andzakovic embraces at once the fetishistic and the spectrophonic angles.

These different hauntograms (and Darija's one in particular) show that a spectrophonic and a reified view on generative AI need not be mutually exclusive. In other words, this technology effectively embodies and performs the Derridean contradiction of automatic autonomy (see Section 1.1). Contradictorily enough (but this is a contradiction of a metaphysics of presence, not of hauntological thinking), AI foregrounds and defers memory and the presence of the remembered, and AI foregrounds and defers its own agency and presence. In doing this, it conceals and at the same time enacts the social relations it appropriates, spectralising itself and the world all around.

These reflections offer enough ground to interpret AI as a *spectral commodity*, a phrase that I clarify as follows. While a table is self-explanatory in its function and use, it nonetheless conceals its "secret" (the social relations from which it derives) quite effectively, since these are nowhere visible in the finished object. AI, by contrast, though opaque in its workings, proves less capable of concealing the presence of the social relationships that constitute it, precisely because it expresses its value by re-enacting them. As *spectrophonia* suggests, these relationships persistently haunt generative AI's outputs in the form of hauntographic (thus, *weird* and *eerie*) traces that betray the human labour from which they derive.

With this reflection, I am not implying that AI might be a transparent technology, but rather a spectral one: a table that visibly stands on its head. As discussed in the introduction to *End of Empire* and as reflected in the installation itself (see Section 5.3), the scale of the appropriation of labour and resources that is performed through AI, in music and well beyond music, is indeed unprecedented. Here, I rather offer a hauntographic perspective to this extractivity, one that points to the *failure of absence* (crucially, not the presence) of the human and non-human bodies and of their social relations that AI algorithmically re-enacts. From the perspective of a metaphysics of presence, this re-enactment is at odds with a "transcendental" individuation of AI (as described by Marx, its acquiring a "wooden brain") through the concealing of its origins. This duality embodies the contradiction of automatic autonomy, a contradiction that is superseded by hauntological thinking.

To extend spectrophonia beyond the musical domain, I describe the traces from the bodies and social relations that the AI commodity appropriates and that foreground its "leaky" nature as *spectrogenic*. Whereas one might say that a text-to-image model generating images

of people with six fingers is hallucinating, the haunting of the accessory finger might be viewed from this perspective as spectrogenic, a spectre that points to the deferred presence (the failure of being present and absent) of the bodies and gestures that a certain model trains upon. Unlike AI hallucinations, spectrogenic traces are not flaws, but rather embody the *weird* and *eerie* spectres of the deferred and deferring ghosts embedded in the data.

7.5.3 Future Work

As a final contribution and further operationalisation of the above reflections, and more as an exploratory and forward-looking move than to provide closure, in this section I extend the hauntographic analysis of the AI commodity through the notion of *exorcism*, which in Chapter 6 I described as a paradoxical act conjuring something away from presence, of declaring the death only to put to death, which inevitably causes the return of the repressed.¹¹³

In *History and Class Consciousness*, Lukács seconds Marx in arguing that commodity fetishism may be viewed as a feature of capitalism whereby “the relation between people takes on the character of a thing and thus acquires a phantom objectivity” (Lukács and Livingstone 1976, 84). According to Lukács’ reading, this phantom objectivity is crucial for capitalist exploitation, as the commodity form comes to mediate all social relations and thus alienates workers from their own labour. Thus, reification naturalises and conceals capitalist extraction and at the same time represents a destabilising factor when raised to consciousness (Lukács and Livingstone 1976, 172).¹¹⁴

Herein lies the connection to AI, whose structural necessity of re-enacting the very relations it seeks to obscure (which the notion of spectrophonia brings to the fore) makes it an inherently unstable and “leaky” commodity. If exposing the reification process destabilises (on a subjective level) the dynamics of alienation, the instability of AI (or at least, of AI in its current stage of evolution) represents a structural breach. And because as we saw the conjuring of spectres is crucial for the AI commodity to perform its function, the question then becomes

113 This definition follows Derrida’s, according to whom exorcisms “convoke the revenant that they conjure away” (Derrida 1994a, 175).

114 While I reserve a comprehensive theoretical grounding of this notion for future work, I employ Lukács because *History and Class Consciousness* proposes a relevant theoretical shift from Marx, in which reification is understood not solely as a consequence of the unplanned relations of capitalist production, but also as a site of (subjective) struggle.

how to hide a ghost in plain sight. Or, perhaps more hauntologically, how to “declare the death only to put to death.” Derrida writes:

But effective exorcism pretends to declare the death only in order to put to death. As a coroner might do, it certifies the death but here it is in order to inflict it. This is a familiar tactic. The constative form tends to reassure. The certification is effective. It wants to be and it must be in effect. It is effectively a performative. But here effectivity phantomalizes itself (Derrida 1994a, 59).

There exist many ways of exorcising a ghost, but “effective” exorcisms are, whether consciously or unconsciously, those wherein death is certified only to be performatively inflicted. Take for instance “Pause Giant AI Experiments: An Open Letter,” the 2023 open letter signed by Elon Musk along with thousands of AI researchers, engineers, and tech leaders, several from the corporations leading LLM development.¹¹⁵ The letter invites AI labs to pause and jointly work on safety protocols, since “powerful AI systems should be developed only once we are confident that their effects will be positive and their risks will be manageable.” Here the engineer-coroner-priest performatively exorcises the very technology he himself helped conjure. The exorcism, though, needs to appeal to well-established and diffuse anxieties to acquire effectivity (the Freudian “Es spukt!” (Derrida 1994a, 166)), anxieties that the engineer-coroner-priest has previously instilled in those attending the performance. Notably, the effect of this exorcism would be the crystallisation of the status quo amid increasing competition, along with one’s position of dominance, and the displacement of more cogent issues into the future by summoning a future threat (I think for instance of ghost labour or data appropriation under “fair use” clauses).

If the above represents a particularly evident form of exorcism, the one I am concerned with here is much more generalised, and refers directly to the dynamics of destabilised reification discussed in Section 7.5. This form of exorcism is exemplified by the so-called “agentic AI” discourse, which through the hauntographic framework I interpret as a tentative exorcism of the spectrogenic traces that haunt AI’s outputs (and threaten the stability of the reification dynamics) by individuating AI as pure agentic form. This, despite research showing that AI agents’ rate of success in the tasks they get assigned is as low as thirty percent (Xu et al. 2025). As further indication of this tendency, I suggest looking no further than what I noted in Section 2.2.2, namely that both trained and untrained AI architectures are commonly

¹¹⁵ https://futureoflife.org/wp-content/uploads/2023/05/FLI_Pause-Giant-AI-Experiments_An-Open-Letter.pdf

described with the word *model*, as if there was no qualitative transition between the two, a tendency in discourse that heavily underplays the relevance of the data.

I thus argue that there exists in the current AI discourse an ongoing conjuration of the spectrogenic traces that haunt generative models:¹¹⁶ by looking at AI in a purely individualised form, it becomes much easier to exorcise the ghost-in-the-room of extraction, appropriation, and commodification of data and labour. To use a metaphor I have been relying on throughout the text, in its individualised form the commodity is placed “at the end of history,” in the Fisherian sense of *out of it* (see Section 6.5.1), severed from the human knowledge it trains upon and rearranges, from its being a (hauntographically unstable) subject of history.

7.6 Conclusions

This chapter concludes the core of the dissertation and precedes a final chapter that summarises the findings. The contributions offered in this chapter stem from six interviews with artists my colleagues at the Intelligent Instruments Lab and I collaborated with in the past three years, which I examined through deductive thematic analysis based on the theoretical probes offered by hauntography. These findings are further supported by my design and artistic practice, and have been developed throughout countless conversations with peers during open labs and daily at the Intelligent Instruments Lab. Through these means, I have sought to answer the third research question that this project asks, relative to *how findings in the intimate human-technology relationship of musical practice aid us in understanding the general impact of AI technologies on our society*. This question, it is worth emphasising, remains an open one, not because the research did not provide enough insight, but because it also opened new research trajectories that I hope to investigate in the years to come.

Based on the hauntographic analysis of the themes emerging from the interviews, I discussed AI NIMes as liminal gateways, *weird thresholds* that separate and at the same time connect worlds, and through which the eerie listening of an alien, machinic outside introduces a weird perspective on empirical reality. Extending Erik Klinga’s comments in Section 7.2.5, AI’s artistic value often lies in how it forces us out of established interpretative schemes, and at the same time compels us to view our world “on its terms.”

¹¹⁶ This argument does not replace, but rather adds to existing commercial reasons and engineering ambitions.

The second theoretical contribution is the notion of *spectrophonia*, which refers to yet another weird disjuncture after the schizophrenic one introduced by Schafer (Schafer 1969). Spectrophonia differs from schizophrenia in the following ways: first, it assumes the schizophrenic act as prior and thus describes a further abstraction from it; second, whereas schizophrenic sounds emerge typically from one localised source, spectrophonic sounds converge from the statistical rearrangements of a multitude; third, while schizophrenia embodies a critique of technology-mediated sonic degradation, spectrophonia focuses on how algorithmically pulverised sounds do or do not bear traces of the multitudes they derive from, investigating strategies to magnify this relationship. Finally, spectrophonia frames AI-generated sounds hauntographically, understanding their sources as hauntographic ghosts: vivifying principles in a state of continuous deferral, whose failure of presence and failure of absence are experienced through spectres in the model's outputs.

The analysis of the interviews, informed by my design and artistic works and supported by conversations with peers in open labs, extends in the latter part into a broader reflection on AI under capitalism. This contribution departs from a hauntographic framing of commodity fetishism. Similarly to all commodities (here, I used Marx's example of a table), AI is weird in that it reifies the social relations that determine it and thus stands upside-down as Marx's table, and *eerie* in that it is prone to assume mystical character and powers. Yet, as the notion of spectrophonia suggests, the social relations that AI appropriates are never fully invisibilised behind the commodity, since they are enacted by the model in use. AI may thus be described as a *spectral commodity* whose "secret" (its own, but also that of any other commodity as described by Marx) is whispered at every prompt through the *weird* and *eerie* remnants of the social relations it appropriates and conjures into and away from presence.

As a springboard for future reflection rather than as a fully developed contribution, I offered an interpretation of the discourse around AI as a form of hauntographic exorcism, aimed at conjuring away the spectrogenic traces that haunt generative models' output. While more work is needed to ground it in material and historical processes, this notion nonetheless provides a new lens by which to observe, for instance, the current "agentic AI" discourse, which I frame hauntographically as an attempt to place AI "at the end of history," severed from the knowledge it trains upon, appropriates, and rearranges. Thus, from its being a subject of history.

8 Conclusions

The structure of this dissertation followed a practice research thread that unravelled throughout the three years of this investigation. The point of departure, aligned with the work package I was entrusted with at the Intelligent Instruments Lab (see Section 1.5), has been design practice of new musical interfaces imbued with AI (Chapter 4), followed by their application in artistic practice (Chapter 5). The findings from practice informed the development of a framework to analyse and create with music AI, and to observe AI as a broader cultural phenomenon (Chapter 6), which I applied systematically through the analysis of interviews stemming from collaborations with artists (Chapter 7).

In this final chapter, I synthesise the findings. Section 8.1 lists the three research questions that guided the investigation, answering each of them. Section 8.2 lists the secondary contributions. I frame these as “secondary” in that they served as a means to answer the research questions rather than answering them directly, but they may arguably be more relevant to specific scholarly niches than the primary contributions advanced in this dissertation (e.g., Thales and Stacco in NIME research, the use of text-to-audio models in performance practice, or the reflections on XAI). Section 8.3 summarises the research journey. Section 8.4 acknowledges some of the limitations in this research, which I view as productive openings for future investigation, and Section 8.5 concludes by outlining three trajectories for future work, all applying hauntographic thinking in different ways.

8.1 Primary Contributions

RQ1: How does AI’s opacity enter and influence the discourse, practices, and aesthetics in artistic creation and performance?

Through practice and user studies, I demonstrated that AI’s resistance to explanation can become, and indeed often represents, a creative resource rather than a technical limitation. Artists often value unpredictability, openness, and emergence as integral to their compositional process, and the opaqueness of generative models contributes to this. This insight contributes to a reorientation of XAI priorities in the arts towards the meaning-making potential of what

remains undisclosed, unfathomed, liminal, which in this dissertation I described with the term Un-XAI.

Whereas indeterminacy was integral to artistic practice long before AI, the more or less unpredictable operations of generative systems foreground a subtle shift. Through artistic practice and by discussing works from other artists, I showed that with generative AI openness and indeterminacy increasingly translate into reflections on liminal presence. Building on Eco, I coined the notion of *haunted work*, wherein liminality becomes a source of poetic meaning, to describe this dynamic.

RQ2: How can insights from these discursive practices contribute to developing tools to interrogate shifting notions of agency, authorship, and intentionality in human-AI co-creation?

The spectral discourse surrounding AI necessitates theoretical tools to query it. Through a comparative reading of Derrida's and Fisher's work on hauntology, I offered the framework of hauntography, which I view as the dissertation's main theoretical contribution. I define hauntography as *a posthumanist framework for the analysis and mobilisation of the traces (the spectres) of what operates by being absent (the ghost) and is apprehended through the affective and fictional categories of the weird and the eerie*.

Hauntography frames the *weird* and the *eerie* as co-constitutive modes of engagement with the spectral, which function both affectively and fictionally, and thus enable analytical and creative applications. This duality grounds a *circle of praxis* in which analytical and creative applications inform each other. While spectral theories often remain elusive in their definitions, hauntography provides a precise vocabulary to investigate and work with spectral phenomena, including notions of *ghost*, *spectre*, *threshold*, *exorcism*, *visor* and *auditor effects*, and a *principle of radical alterity* that opposes anthropocentric projection. Section 6.2 provides the complete hauntographic glossary.

RQ3: How do findings in the intimate human-technology relationship of musical practice aid us in understanding the general impact of AI technologies on our society?

Through thematic analysis based on the hauntographic framework of interviews with artists I collaborated with during the research, I offered three insights into music AI and AI's broader

cultural and political dimensions, which I structure below as distinct but interrelated contributions.

First, a reflection on AI NIMEs as hauntographic *thresholds*. Drawing on Fisher's distinction between *fantasy fiction*, where thresholds are conduits into worlds entirely detached from ours, and *weird fiction*, where thresholds introduce a *weird* outside to destabilise the structures of the inside, I described AI NIMEs as *weird thresholds*, which introduce a non-human perspective on the training data and the things in the world it signifies.

Second, the notion of *spectrophonia*. Paralleling R. Murray Schafer's notion of schizophonia, spectrophonia describes how in generative AI algorithmically pulverised schizophonic sounds converge from the dataset's multitude, bearing spectral traces of the sources understood as hauntographic ghosts (and thus always already deferred). While schizophonic sounds typically maintain a certain aura of authenticity, in the case of spectrophonia this process is further complicated, in that the generative model contests the aura of the sources it ruminates and rearranges. One's subjective relationship with the training data counterbalances this pull by magnifying the spectrophonic process.

Third, I identified another hauntographic dynamic in the reification of the model that individuates it outside of the social relations constituting the training data. I analysed the interplay between the two hauntographic processes identified (spectrophonia on the one hand, the model's reification on the other) through the notion of *commodity fetishism*. Whereas commodities typically succeed in concealing the social relations from which they derive, AI must re-enact these very relations. This makes it an inherently unstable, spectral commodity, haunted by the *failure of absence* of the social relations it appropriates. This instability may be further analysed through the hauntographic notion of exorcism, interpreting the current discourse around AI as an attempt to complete the reification by individuating AI as pure agentic form. Thus, AI's current status as a spectral commodity positions it as a privileged site for understanding and revealing broader structures of capitalist extraction. This represents a research trajectory I intend to develop in future work.

8.2 Secondary Contributions

As the research progressed, a series of secondary contributions has emerged. In this section, I list and summarise some key secondary contributions, dividing them into three groups: design

of AI interfaces, artistic works, theoretical contributions. Within each cluster, contributions are ordered based on their appearance in the dissertation and referenced in the body of the text.

8.2.1 AI NIMES

Thales: A NIME for latent space navigation in NAS. Thales uses reciprocal magnetic repulsion between two hand-held controllers, or between the controllers and magnetically inscribed surfaces, to produce the effect of a tangible yet invisible object. *See Section 4.3.*

Stacco: A NIME for latent space navigation in NAS. It enables real-time navigation of NAS latent space through the displacement of magnetic marbles around an oval wooden board. Stacco supports notation through *embodied sketching*, and has undergone three major iterations to address different artistic and technical needs. *See Section 4.4.*

Embodied Sketching: A notational technique adopted to design scores with NAS latent navigation. The idea is to apply oval cardboard sheets and tracing paper on Stacco's surface. *See Section 4.4.3.*

Mouja+ Pipeline: A performative system combining asynchronous text-to-audio generation with real-time NAS for timbre transfer and latent manipulation. The pipeline leverages the complementary aspects of text-to-audio and real-time NAS, the former being asynchronous but enabling macro formal control, the latter being synchronous but allowing mainly timbral manipulation. *See Section 5.4.2.*

8.2.2 Artistic Works

Mouja: A performance with Stacco and Thales based on séance aesthetics. *See Section 5.2.*

End of Empire: A large-scale participatory, site-specific installation with textiles and NAS. *See Section 5.3.*

Mouja+: A performance piece featuring Thales and Stacco in an Ambisonics dome, which extends the practice of prompt jockeying with NAS to live use by combining text-to-audio and real-time NAS. *See Section 5.4.*

8.2.3 Theoretical Contributions

Performative and Compositional Challenges with NAS: I identified and analysed three compositional and performative challenges with NAS that emerged in my practice and in working with other artists. These are *arbitrary encoding*, *feature entanglement*, and *dimensional entanglement*. *See Section 4.2.2.*

Magnetic Scores: A notion describing inherent scores in which the musical inscription emerges relationally through magnetic interactions. *See Section 4.2.3.*

Embodied Understanding of Latent Topologies: A study approaching XAI from a pragmatic and embodied angle. The study features Stacco and a series of NAS adapters developed by Victor Shepardson. The results indicate that the unpredictability and surprise of an unregularised latent space can be positively valued by artists. *See Section 4.6.3.*

Practices Towards a Non-extractivist Use of AI: In working on *End of Empire*, I demonstrated a participatory, non-extractivist approach to the use of AI technologies in artistic practice. The approach involved community engagement in all phases, including data collection through sound walks and individual listening sessions, the use of small-data architectures, and their local deployment. *See Section 5.3.3.*

Compositional Strategies with text-to-audio NAS: I offered three strategies for designing prompts for text-to-audio models. These are *dataset exploration*, *strategic use of negative prompts*, *prompt micro-tuning*. *See Section 5.4.5.*

Dataset Accessibility in Artistic Practice: I demonstrated how open access to training data and effective search tools improve the engagement with generative AI within and potentially beyond artistic practice. *See Section 5.4.6.*

8.3 Research Journey

The arc of this research unfolded through interface design (Chapter 4), artistic practice (Chapter 5), cultural theory (Chapter 6), and analysis (Chapter 7). I departed from the notion of inherent scores (open systems in which the musical interface and the notation merge with unprecedented fluidity) and discussed their relevance to NAS. I then introduced the term *magnetic scores*. The latter provides a theoretical umbrella for Thales and Stacco, the two NIME design contributions that the dissertation offers.

Through observations from performance practice with Thales I argued that the sustained engagement with a certain AI system can become in itself an explanatory strategy. Based on results from a user study with Stacco, I argued that the resistance to explanation that XAI seeks to resolve can, under certain conditions, become a productive constraint. These insights motivate the research focus on the meaning-making potential of AI's opacity.

Chapter 5 expanded on this intuition through three artistic contributions: *Mouja*, investigating presence through séance aesthetics; *End of Empire*, exploring a participatory approach to dataset collection and the estrangement of community memory through NAS; *Mouja+*, reflecting on the use of natural language with a generative model, on the relevance of dataset accessibility, and on the relationship of textual and sonic signifiers to referents that never were. The analysis of these artistic works, sustained by observations on works by other artists, led to the notion of *haunted work*, reflecting on how indeterminacy in AI art often foregrounds a question of presence.

To further develop and operationalise this notion, Chapter 6 offered a comparative analysis of Derrida's *Specters of Marx* and Fisher's broader corpus of work, focusing in particular on *The Weird and the Eerie*. The hauntographic framework offers a vocabulary (outlined in Section 6.2) to investigate phenomena and create through the lens of spectrality.

Chapter 7 applied hauntography through deductive thematic analysis of interviews with six artists. Through this analysis, I discussed AI NIMES as *weird thresholds*, introduced the notion of *spectrophonia*, and drawing on Marx described AI as a *spectral commodity*, structurally incapable of fully concealing the social relations it reifies (see Section 7.5).

8.4 Limitations

Before addressing the question of how I envision this research to progress beyond these pages, it is necessary to reflect on some of its limitations. I divide these into two clusters, one concerning structural limitations of the work, and one concerning more localised limitations, often relative to the methods or cohorts involved in specific activities. I address these chapter by chapter.

Among the first cluster, of particular relevance are three limitations. First amongst these is the wide scope of the dissertation, spanning XAI questions, instrument design, artistic practice and cultural studies. This breadth inevitably follows from the Intelligent Instruments Lab's research orientation, investigating how AI is shaping society through music practice. This required developing the musical instruments that would embed and activate the AI algorithms (Chapter 4), developing a practice with them (Chapter 5), abstracting from the practice new theoretical insights (Chapter 6), and extending these insights beyond artistic practice in order to gain new knowledge (Chapter 7).

Second, in practice research the very questions informing one's investigation, and not just the answers, often become apparent as practice progresses. Thus, there exists a messy and emergent character which also entails a certain unpredictability of the research trajectory. In other words, new routes and dead ends are part of the journey, and require an attitude of openness and trust in the process. This inevitably entailed the opening of threads that I could not fully pursue (e.g. the relationship between Barad's agential realism and Derrida's hauntology, or the final excursion on AI as a cultural exorcism), in that each would deserve a research project of its own.

A third limitation is the dissertation's limited engagement with non-Western frameworks. In future work I would like to further reflect on hauntography through other systems of thinking (I think for instance of Afrofuturism, which I touched upon in Chapter 6, and which shares some crucial aspects with Fisher's framing of hauntology).

The second cluster of limitations I would like to address is instead local, and thus deserves to be contextualised chapter by chapter.

Chapter 4 presents two studies dealing with XAI from a similarly pragmatic and embodied angle, but through different methods (the former, through practice research, the latter through a qualitative study followed by interviews). Both studies provided insights that redefined the research trajectory. While I consider these sufficient to motivate the shift towards

Un-XAI described in the chapter, the specific contributions they bring deserve to be further validated with additional empirical research.

Chapter 5 presents three performances that I discuss from the perspective of spectrality as defined in Section 2.5, before introducing hauntography in Chapter 6. An alternative option in structuring the dissertation would have been to open with hauntography, and then analyse the performances hauntographically. I reflected extensively on this possibility, and concluded that while the latter approach would have provided closure to the circle of praxis I discuss in Chapter 6, it would not have accounted for *how* hauntography emerged from practice itself. I nonetheless consider a practice-driven, systematic application of hauntography as a crucial research trajectory to pursue in future work.

Finally, in Chapter 7 I present interviews with six artists I collaborated with in the past three years. I consider this number sufficient to begin articulating the findings for two reasons: first, because each interview is supported by months of sustained collaboration with the artist; second, because the reflections I offered are supported by countless discussions with peers during open labs, reflections with students, and my own artistic and design practice. Nonetheless, in future work I intend to provide further empirical scaffolding to the notions developed in this chapter.

8.5 Future Work

Throughout the arc of this dissertation, there has been no lack of cues for new research trajectories. Among these, three resonate in particular with how I envisage myself as a scholar and as an artist in the years to come. Of these, two are related to AI, and one hints instead at the application of hauntography for a different type of research.

Hauntography in Artistic Practice: The dissertation concluded with the application of hauntography as an analytical tool through the interviews described in Chapter 7. Yet, as outlined in Section 6.6.2, hauntography inherits Fisher's framing of the *weird* and the *eerie* as both affective modes and modes of fiction. I would therefore like to give a subjective closure to the hauntographic circle of praxis by bringing the analytical categories I developed in Chapter 7 back into my own artistic practice.

In particular, I consider the notion of spectrophonia promising, certainly as a theoretical construct per se, but also in its practical application in the arts, since it reflects on how the subjective understanding of and engagement with the sonic material shapes the understanding of a certain AI system. In future work, I thus intend to investigate how the principles I briefly described in Section 7.4.2, but also the extensive corpus of artistic practices focusing on the modern soundscape, might inform the artistic engagement with the contemporary AI soundscape.

Hauntography in AI Culture: While Chapter 7 already offered an analytical hauntography, the insights that derived from it deserve to be further tested and refined against a much broader sample of use cases. As I perfected the theoretical framework offered in Chapter 6, I began applying it quite intuitively to the cultural, political, and mediatic landscape that surrounds me. This lens considerably refined my understanding of cultural phenomena. As a second research trajectory, which as per hauntography's circle of praxis I do not consider parallel to but rather intertwined with artistic and design practice, I would like to further analyse hauntographically the economic, cultural and political structures from which the phenomenon we describe as AI emerges. This is, of course, an extensive and challenging task, but the instruments to attempt it have perhaps by now matured enough.

Hauntography and the Posthuman: In Chapter 6, I discussed Fisher's critical position toward the anthropomorphising move performed by some posthumanist theories. Hauntography's principle of radical alterity derives precisely from this critique. I consider Fisher's assessment of new materialisms in *The Weird and the Eerie* slightly hurried, but I do share his concern as far as this relates to the way posthumanist theories oftentimes translate discursively. If circumstances permit, in future critical work I would like to address how posthumanist theories shift, transform, and at times distort as they enter the public sphere. The theoretical and methodological resources for this investigation are already established in the framework introduced in this dissertation and the participatory methods I learned and contributed to at the Intelligent Instruments Lab.

8.6 Epilogue

Stepping back from the research questions and adopting a bird's-eye perspective on the dissertation as a whole, I am inclined to think that its most transformative aspect might lie less in the answers it provides than in the premises it adopts. Using AI's opacity as a creative and epistemic resource has proven a rewarding approach, opening an unexplored research path rich in theoretical and creative insight.

Inevitably, this epistemological positioning has also shaped how I understand the world around me, and how I have come to think of myself as haunted. Indeed, as an artist working with generative systems at a moment of sustained reflection on what it means to decentre the human, in writing these pages I was constantly confronted with the tacit question of what a posthumanism of ghosts might look like, within and beyond AI. As hauntography suggests, it might be one in which the contradiction of automatic autonomy comes to haunt our own autonomy and automatisms, unsettling how we understand ourselves as separate from the material and mundane contingencies that surround us: always already at the end of history, even as we begin anew.

There now exists a technology that does what the Voder did, but with greater fidelity and realism. One haunted by thousands of Helen Harpers. The spectrophonic traces we have encountered within and all around its hidden layers have made it perhaps all too evident that “we are they, and they are we.”¹¹⁷

¹¹⁷ Fisher 2016, 109.

9 References

- Agamben, G. (2000) *Means without End: Notes on Politics*. Minneapolis: University of Minnesota Press.
- Andelic-Andzakovic, D., Duarte, N. and Privato, N. (2023) ‘The historically informed AI: Johann Sperger through time and bass’, *AIMC 2023*.
- Armitage, J., Privato, N., Shepardson, V. and Betancur Gutierrez, C. (2023) ‘Explainable AI in music performance: Case studies from live coding and sound spatialisation’, *XAI in Action: Past, Present, and Future Applications, NeurIPS 2023 Workshop*. New Orleans, USA. Available at: <https://openreview.net/forum?id=se4ojQqjB>.
- Armstrong, J. (2010) ‘On the possibility of spectral ethnography’, *Cultural Studies ↔ Critical Methodologies*, 10(3), pp. 243–250. doi: 10.1177/1532708609359510.
- Arora, C. and Sarkar, D. (2023) ‘On the injunction of XAIxArt’, arXiv preprint. doi: 10.48550/arXiv.2309.06227.
- Austin, J.L. (1975) *How to Do Things with Words*. Cambridge: Harvard University Press.
- Barad, K. (2003) ‘Posthumanist performativity: Toward an understanding of how matter comes to matter’, *Signs: Journal of Women in Culture and Society*, 28(3), pp. 801–831. doi: 10.1086/345321.
- Barad, K. (2014) ‘Diffracting diffraction: Cutting together-apart’, *Parallax*, 20(3), pp. 168–187. doi: 10.1080/13534645.2014.927623.
- Blanco, M.D. and Peeren, E. (2013) *The Spectralities Reader: Ghosts and Haunting in Contemporary Cultural Theory*. London: Bloomsbury Publishing.
- Boden, M.A. (2014) ‘GOF AI’, in Frankish, K. and Ramsey, W.M. (eds.) *The Cambridge Handbook of Artificial Intelligence*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9781139046855.007.
- Bogost, I. (2012) *Alien Phenomenology, or What It’s Like to Be a Thing*. Minneapolis: University of Minnesota Press.

- Bolt, B. and Barrett, E. (eds.) (2019) *Practice as Research: Approaches to Creative Arts Enquiry*. London: I.B. Tauris. doi: 10.5040/9780755604104.
- Born, G. and Barry, A. (2018) 'Music, mediation theories and actor-network theory', *Contemporary Music Review*, 37(5–6), pp. 443–487. doi: 10.1080/07494467.2018.1578107.
- Boulez, P. (2001) 'Letter to Karlheinz Stockhausen, dated 25 September 1957', in Misch, I. and Bandur, M. (eds.) *Karlheinz Stockhausen bei den Internationalen Ferienkursen für Neue Musik in Darmstadt 1951–1996: Dokumente und Briefe*. Kürten: Stockhausen-Verlag, pp. 179–181.
- Bourriaud, N. (2009) *Relational Aesthetics*. Dijon: Les Presses du Réel.
- Brey, P. (2000) 'Technology and embodiment in Ihde and Merleau-Ponty', in Mitcham, C. (ed.) *Metaphysics, Epistemology, and Technology*. Research in Philosophy and Technology, vol. 19. JAI Press, pp. 45–58.
- Briot, J.-P., Hadjeres, G. and Pachet, F.-D. (2017) 'Deep learning techniques for music generation – a survey', arXiv preprint. doi: 10.48550/arXiv.1709.01620.
- Brown, C. (2013) 'What is “comparative” literature?', *Comparative Critical Studies*, 10(1), pp. 67–88. doi: 10.3366/ccs.2013.0077.
- Bryan-Kinns, N., Banar, B., Ford, C. et al. (2023a) 'Exploring XAI for the arts: Explaining latent space in generative music', arXiv preprint. doi: 10.48550/arXiv.2308.05496.
- Bryan-Kinns, N., Ford, C., Chamberlain, A. et al. (2023b) 'Proceedings of the first international workshop on eXplainable AI for the arts (XAIxArts)', arXiv preprint. doi: 10.48550/arXiv.2310.06428.
- Bryan-Kinns, N., Wszeborowska, A., Sutskova, O., Wilson, E. et al. (2025) 'Leveraging small datasets for ethical and responsible AI music making', *Audio Mostly 2025*. doi: 10.1145/3771594.3771601.
- Bulley, J. and Şahin, O. (2021) *Practice Research – Report 1: What Is Practice Research? and Report 2: How Can Practice Research Be Shared?* London: British Library. doi: 10.23636/1347.

- Caillon, A. and Esling, P. (2021) 'RAVE: A variational autoencoder for fast and high-quality neural audio synthesis', arXiv preprint. doi: 10.48550/arXiv.2111.05011.
- Callon, M. and Latour, B. (1992) 'Don't throw the baby out with the bath school! A reply to Collins and Yearley', in Pickering, A. (ed.) *Science as Practice and Culture*. Chicago: University of Chicago Press, pp. 343–368.
- Calvino, I. (1987) *The Uses of Literature*. New York: Harper Collins Publishers.
- Candy, L. (2006) *Practice-based Research: A Guide*. Creativity and Cognition Studios Report 2006-V1.0. Sydney: Creativity and Cognition Studios, University of Technology.
- Carvalho, D.V., Pereira, E.M. and Cardoso, J.S. (2019) 'Machine learning interpretability: A survey on methods and metrics', *Electronics*, 8(8), p. 832. doi: 10.3390/electronics8080832.
- Clarke, A.C. (1958) *Voice Across the Sea*. London: Frederick Muller.
- Clarke, M. (2025) 'Sonic experimentalism in the practice of conducting: Sound sticks and magic wands, or expansion of practice?', *Leonardo*, pp. 337–343. doi: 10.1162/leon_a_02706.
- Cristianini, N. (2021) 'Shortcuts to artificial intelligence', in Pelillo, M. and Scantamburlo, T. (eds.) *Machines We Trust*. Cambridge: The MIT Press. doi: 10.7551/mitpress/12186.003.0006.
- Dahlstedt, P. (2025) 'Adaptive elusion – an improvisation for pianist and real-time machine learning', *International Conference on New Interfaces for Musical Expression*. doi: 5281/zenodo.17801166.
- Dang, Y., Huang, K., Huo, J. et al. (2024) 'Explainable and interpretable multimodal large language models: A comprehensive survey', arXiv preprint. doi: 10.48550/arXiv.2412.02104.
- De Souza, J. (2017) *Music at Hand: Instruments, Bodies, and Cognition*. Oxford: Oxford University Press.

- Deleuze, G. and Guattari, F. (1983) *Anti-Oedipus: Capitalism and Schizophrenia*. Minneapolis: University of Minnesota Press.
- Denzin, N.K. and Lincoln, Y.S. (eds.) (2003) *The Landscape of Qualitative Research: Theories and Issues*. Los Angeles: Sage.
- Derrida, J. (1994a) *Specters of Marx: The State of the Debt, the Work of Mourning, and the New International*. London: Routledge.
- Derrida, J. (1994b) 'Spectres of Marx', *New Left Review*, I(205), pp. 31–58. doi:10.64590/jrf.
- Devis, N., Demerlé, N., Nabi, S., Genova, D. and Esling, P. (2023) 'Continuous descriptor-based control for deep audio synthesis', arXiv preprint. doi: 10.48550/arXiv.2302.13542.
- Dick, P.K. (1984) *Time out of Joint*. Bournemouth: Bluejay Books.
- Dickens, C. (1847) *Dombey and Son*. London: Bradbury and Evans.
- Donnarumma, M. (2011) 'Xth Sense: Researching biological sounds of human body for an experimental paradigm of musical performance', *Proceedings of the International Computer Music Conference (ICMC)*, Huddersfield.
- Dudley, H.W. (1939) *VODER* [Demonstration film]. Bell Telephone Laboratories. Available at: <https://archive.org/details/voder1939>.
- Eco, U. (1989) *The Open Work*. Cambridge: Harvard University Press.
- Einarsson, E.T. (2017) 'Re-notations: Flattening hierarchies and transforming functions', in *The Dark Precursor: Deleuze and Artistic Research*. Leuven: Leuven University Press.
- Engel, J., Resnick, C., Roberts, A. et al. (2017) 'Neural audio synthesis of musical notes with WaveNet autoencoders', arXiv preprint. doi: 10.48550/arXiv.1704.01279.
- Evans, Z., Parker, J.D., Carr, C.J., Zukowski, Z., Taylor, J. and Pons, J. (2024) 'Stable Audio Open', arXiv preprint. doi: 10.48550/arXiv.2407.14358.
- Feld, S. (1996) 'Pygmy POP. A genealogy of schizophonic mimesis', *Yearbook for Traditional Music*, 28, pp. 1–35. doi: 10.2307/767805.

- Fiebrink, R., Trueman, D. and Cook, P.R. (2009) 'A meta-instrument for interactive, on-the-fly machine learning', *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.1177513.
- Fisher, M. (2010) *Capitalist Realism: Is There No Alternative?* Winchester: Zero Books.
- Fisher, M. (2013a) *Ghosts of My Life: Writings on Depression, Hauntology and Lost Futures*. Alresford: Zero Books.
- Fisher, M. (2013b) 'The metaphysics of crackle: Afrofuturism and hauntology', *Dancecult*, 5(2), pp. 42–55. doi: 10.12801/1947-5403.2013.05.02.03.
- Fisher, M. (2016) *The Weird and the Eerie*. London: Repeater Books.
- Freud, S. (1919) 'The Uncanny', *Imago*, Bd. V; reprinted in *Sammlung, Fünfte Folge*. Translated by Alix Strachey.
- Friedman, K. and Smith, O. (2002) *Fluxus Performance Workbook*. London: Performance Research e-publication.
- Fukuyama, F. (1989) 'The end of history?', *The National Interest*, (16), pp. 3–18. Available at: <https://www.jstor.org/stable/24027184>.
- Fukuyama, F. (2006) *The End of History and the Last Man*. New York: Free Press.
- Gioti, A. (2021) 'A Compositional Exploration of Computational Aesthetic Evaluation and AI Bias'. *International Conference on New Interfaces for Musical Expression*. doi: 10.21428/92fbeb44.de74b046.
- Grant, C. and Osanloo, A.F. (2014) 'Understanding, selecting, and integrating a theoretical framework in dissertation research: creating the blueprint for your "house"', *Administrative Issues Journal*, 4, pp. 12–26. doi: 10.5929/2014.4.2.9.
- Gray, C. (1996) *Inquiry through Practice: Developing Appropriate Research Strategies*. Available at: <http://www.carolegray.net/Papers%20PDFs/ngnm.pdf>.
- Haas, P.M. (1992) 'Introduction: Epistemic communities and international policy coordination', *International Organization*, 46(1), pp. 1–35. doi: 10.1017/S0020818300001442.

- Hägglund, M. (2008) *Radical Atheism: Derrida and the Time of Life*. Stanford: Stanford University Press.
- Hamer, F.T.C. (2023) *Developing Memory: Remembrance, Embellishment and Hauntography*. PhD thesis. Montreal: Concordia University.
- Hao, K. (2025) *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*. New York: Penguin Press.
- Haraway, D. (1988) 'Situated knowledges: The science question in feminism and the privilege of partial perspective', *Feminist Studies*, 14(3), pp. 575–599. doi: 10.2307/3178066.
- Hardinges, J., Simperl, E. and Shadbolt, N. (2023) 'We must fix the lack of transparency around the data used to train foundation models', *Harvard Data Science Review*, Special Issue 5. doi: 10.1162/99608f92.a50ec6e6.
- Haseman, B. (2006) 'A manifesto for performative research', *Media International Australia*, 118(1), pp. 98–106. doi: 10.1177/1329878X0611800113.
- Hawas, M. (2025) "'Aren't they all in the literary histories?": English criticism and the dismissal of comparative literature', *European Review*, 33(3), pp. 360–370. doi: 10.1017/S1062798725100203.
- Hegel, G.W.F. (2010) *Aesthetics: Lectures on Fine Art*. Oxford: Clarendon Press.
- Herremans, D., Chuan, C.-H. and Chew, E. (2018) 'A functional taxonomy of music generation systems', *ACM Computing Surveys*, 50(5), pp. 1–30. doi: 10.1145/3108242.
- Hui, Y. (2023) 'ChatGPT, or the eschatology of machines', *E-Flux*, (137). Available at: <https://www.e-flux.com/journal/137/544816/chatgpt-or-the-eschatology-of-machines/>.
- Ihde, D. (1996) *Technology and the Lifeworld: From Garden to Earth*. Bloomington: Indiana University Press.
- Impett, J. (2023) 'Music, discourse and intuitive technology', *AI & Society*, 38(5), pp. 1885–1896. doi: 10.1007/s00146-020-01126-4.
- Jameson, F. (1991) *Postmodernism, or, The Cultural Logic of Late Capitalism*. Durham: Duke University Press.

- Juzek, T.S. and Ward, Z.B. (2024) ‘Why does ChatGPT “delve” so much? Exploring the sources of lexical overrepresentation in large language models’, arXiv preprint. doi: 10.48550/arXiv.2412.11385.
- Kingma, D.P. and Welling, M. (2019) *An Introduction to Variational Autoencoders*. doi: 10.48550/arXiv.1906.02691.
- Kittler, F.A. (1990) *Discourse Networks, 1800/1900*. Stanford: Stanford University Press.
- Kivunja, C. (2018) ‘Distinguishing between theory, theoretical framework, and conceptual framework: A systematic review of lessons from the field’, *International Journal of Higher Education*, 7(6), p. 44. doi: 10.5430/ijhe.v7n6p44.
- Lakoff, G. and Johnson, M. (1980) ‘The metaphorical structure of the human conceptual system’, *Cognitive Science*, 4(2), pp. 195–208. doi: 10.1207/s15516709cog0402_4.
- Latour, B. (2005) *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press.
- Law, J. (2023) ‘Material semiotics’, in Atkinson, P., Delamont, S., Cernat, A., Sakshaug, J.W. and Williams, R.A. (eds.) *SAGE Research Methods Foundations*. London: SAGE. doi: 10.4135/9781529751079747550.
- Lepri, G. (2022) *Play-Make-Believe: Discovering the Musical Values Inscribed into Digital Instruments*. PhD thesis. Queen Mary University of London. Available at: <http://instrumentslab.org/data/giacomo/LepriPhDThesis.pdf>.
- Lepri, G., Privato, N. and Magnusson, T. (2024) ‘Embodied sketching for neural synthesis’, *Audio Mostly 2024 – Explorations in Sonic Cultures*, pp. 549–551. doi: 10.1145/3678299.3678358.
- Lewis, P., Perez, E., Piktus, A. et al. (2020) ‘Retrieval-augmented generation for knowledge-intensive NLP tasks’, arXiv preprint. doi: 10.48550/arXiv.2005.11401.
- Linardatos, P., Papastefanopoulos, V. and Kotsiantis, S. (2020) ‘Explainable AI: A review of machine learning interpretability methods’, *Entropy*, 23(1), p. 18. doi: 10.3390/e23010018.

- Lynch, D. (2001) *Mulholland Drive* [Film]. Los Angeles: Universal Pictures.
- Lipton, Z.C. (2016) 'The mythos of model interpretability', arXiv preprint. doi: 10.48550/arXiv.1606.03490.
- Longpre, S., Mahari, R., Lee, A. et al. (2024) 'Consent in crisis: The rapid decline of the AI data commons', arXiv preprint. doi: 10.48550/arXiv.2407.14933.
- Losurdo, D. (2017) *Il marxismo occidentale: come nacque, come morì, come può rinascere*. Bari: GLF editori Laterza.
- Lucier, A. (1998) 'Origins of a form: Acoustical exploration, science and incessancy', *Leonardo Music Journal*, 8, pp. 5–11. doi: 10.2307/1513391.
- Lukács, G. (1976) *History and Class Consciousness: Studies in Marxist Dialectics*. Translated by R. Livingstone. Cambridge: MIT Press.
- Luque, S. (2009) 'The stochastic synthesis of Iannis Xenakis', *Leonardo Music Journal*, 19, pp. 77–84. doi: 10.1162/lmj.2009.19.77.
- Lyotard, J.-F. (1984) *The Postmodern Condition: A Report on Knowledge*. Minneapolis: University of Minnesota Press.
- Magnusson, T. (2009) 'Of epistemic tools: Musical instruments as cognitive extensions', *Organised Sound*, 14(2), pp. 168–176. doi: 10.1017/S1355771809000272.
- Magnusson, T. (2010) 'Designing constraints: Composing and performing with digital musical systems', *Computer Music Journal*, 34(4), pp. 62–73.
- Magnusson, T. (2011) 'Algorithms as scores: Coding live music', *Leonardo Music Journal*, 21, pp. 19–23. doi: 10.1162/LMJ_a_00056.
- Magnusson, T. (2019) *Sonic Writing: Technologies of Material, Symbolic and Signal Inscriptions*. London: Bloomsbury Academic.
- Magnusson, T., Braybrooke, K., Skach, S. et al. (2024) 'The Organium: A library of technical elements for improvisatory design thinking', *Proceedings of the 19th International Audio Mostly Conference: Explorations in Sonic Cultures*, pp. 198–208. doi: 10.1145/3678299.3678319.

- Mannone, M.C., Kitamura, E., Huang, J., Sugawara, R. and Kitamura, Y. (2018) 'Cubeharmonic: A new interface from a magnetic 3D motion tracking system to music performance', *Proceedings of the International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.1302615.
- Marx, K. (1887/1995) *Capital: A Critique of Political Economy. Volume I*. Translated by S. Moore and E. Aveling. Edited by F. Engels. Moscow: Progress Publishers.
- Marx, K. and Engels, F. (1979) *Manifesto of the Communist Party*. New York: International Publishers.
- Massumi, B. (2002) *Parables for the Virtual: Movement, Affect, Sensation*. Durham: Duke University Press. doi: 10.1215/9780822383574.
- McPherson, A., Jack, R. and Moro, G. (2016) 'Action-sound latency: Are our tools fast enough?', *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.3964611.
- Merendino, N., Lepri, G., Rodà, A. and Masu, R. (2023) 'Redesigning the Chowndolo: A reflection-on-action analysis to identify sustainable strategies for NIMEs design', *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.11189153.
- Merleau-Ponty, M. (1990) *Phenomenology of Perception*. Abingdon-on-Thames: Taylor & Francis.
- Miller, T. (2019) 'Explanation in artificial intelligence: Insights from the social sciences', *Artificial Intelligence*, 267, pp. 1–38. doi: 10.1016/j.artint.2018.07.007.
- Mitterhuber, M., Tomás, E. and Kaltenbrunner, M. (2025) 'Ottosonics: Designing and implementing an open and affordable 3D spatial audio system', *Organised Sound*, 2025;30(2):191-203. doi: 10.1017/S1355771825100563.
- Möck, L.A. (2022) 'Prediction promises: Towards a metaphorology of artificial intelligence', *Journal of Aesthetics and Phenomenology*, 9(2), pp. 119–139. doi: 10.1080/20539320.2022.2143654.

- Morrison, L. and McPherson, A. (2024) ‘Entangling entanglement: A diffractive dialogue on HCI and musical interactions’, *CHI Conference on Human Factors in Computing Systems*, pp. 1–17. doi: 10.1145/3613904.3642171.
- Morton, T. (2013) *Hyperobjects: Philosophy and Ecology after the End of the World*. Minneapolis: University of Minnesota Press.
- Nyman, M. (1974) *Experimental Music: Cage and Beyond*. Cambridge: Cambridge University Press.
- Nyrup, R. and Robinson, D. (2022) ‘Explanatory pragmatism: A context-sensitive framework for explainable medical AI’, *Ethics and Information Technology*, 24(1), p. 13. doi: 10.1007/s10676-022-09632-3.
- Oliveros, P. (2013) *Anthology of Text Scores*. Edited by S. Golter and L. Hall. Kingston, New York: Deep Listening Publications.
- Oliveros, P. (1984) *Software for People: Collected Writings 1963–80*. Baltimore: Smith Publications.
- Packham, J. (2024) ‘Séance and technology: Intermundane communication as a methodology for contemporary music’, *International Journal of Performance Arts and Digital Media*, 20(3), pp. 464–483. doi: 10.1080/14794713.2024.2412507.
- Páez, A. (2020) *The Pragmatic Turn in Explainable Artificial Intelligence (XAI)*. arXiv preprint. doi: 10.48550/arXiv.2002.09595.
- Paradiso, J.A., Hsiao, K. and Benbasat, A. (2001) ‘Tangible music interfaces using passive magnetic tags’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.1176374.
- Pardue, L.S. and Paradiso, J.A. (2002) ‘Musical Navigatrics: New musical interactions with passive magnetic tags’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.1176456.
- Pasquinelli, M. (2023) *The Eye of the Master: A Social History of Artificial Intelligence*. London: Verso.

- Pelinski, T., Diaz, R., Benito Temprano, A.L. and McPherson, A. (2023) ‘Pipeline for recording datasets and running neural networks on the Bela embedded hardware platform’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.11189141.
- Pickering, A. (1995) *The Mangle of Practice: Time, Agency, and Science*. Chicago: University of Chicago Press.
- Polanyi, M. (1966) *The Tacit Dimension*. New York: Doubleday.
- Potochnik, A. (2016) ‘Scientific explanation: Putting communication first’, *Philosophy of Science*, 83(5), pp. 721–732. doi: 10.1086/687858.
- Privato, N. and Armitage, J. (2023) ‘A context-sensitive approach to XAI in music performance’, arXiv preprint. doi: 10.48550/arXiv.2309.04491.
- Privato, N., Lepri, G., Magnusson, T. and Einarsson, E.T. (2024a) ‘Sketching magnetic interactions for neural synthesis’, *International Conference on Technologies for Music Notation and Representation*. Zurich: Zurich University of the Arts.
- Privato, N. and Novello, A. (2021) ‘Generative scores and data mining: W.E.I.R.D. enters the stage’, *Proceedings of the 9th xCoAx – Conference on Computation, Communication, Aesthetics & X*.
- Privato, N., Rampado, O. and Novello, A. (2022) ‘A creative tool for the musician combining LSTM and Markov chains in MaxMSP’, *EvoMUSART 2022*. Berlin: Springer, pp. 228–242. doi: 10.1007/978-3-031-03789-4_15.
- Privato, N., Magnusson, T. and Einarsson, E.T. (2023a) ‘Magnetic interactions as a somatosensory interface’, *International Conference on New Interfaces for Musical Expression*. Mexico City. doi: 10.5281/zenodo.11189218.
- Privato, N. (2024b) ‘Mouja’, *12th Conference on Computation, Communication, Aesthetics & X*. Available at: <https://2024.xcoax.org/pdf/privato.pdf>.
- Privato, N. and Magnusson, T. (2025) ‘The open work and AI: Co-creation and haunting in text scores and inherent scores’. *Organised Sound*. 30(3): pp. 290-300. doi: 10.1017/S135577182510085X.

- Privato, N., Magnusson, T. and Einarsson, E.T. (2023b) ‘The magnetic score: Somatosensory inscriptions and relational design in the instrument-score’, *International Conference on Technologies for Music Notation and Representation (TENOR 2023)*. Boston: Northeastern University, pp. 36–44.
- Privato, N., Shepardson, V., Lepri, G. and Magnusson, T. (2024c) ‘Stacco: Exploring the embodied perception of latent representations in neural synthesis’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.13904899.
- Raffel, C. (2016) *Learning-Based Methods for Comparing Sequences, with Applications to Audio-to-MIDI Alignment and Matching*. PhD thesis. Columbia University. doi: 10.7916/D8N58MHV.
- Raven, J.E. and Kirk, G.S. (1957) ‘Chapter II: Thales of Miletus’, in *The Presocratic Philosophers*. Cambridge: Cambridge University Press, pp. 74–98.
- Raza, S., Ghuge, S., Ding, C., Dolatabadi, E. and Pandya, D. (2024) ‘FAIR enough: Develop and assess a FAIR-compliant dataset for large language model training?’, *Data Intelligence*, 6(2), pp. 559–585. doi: 10.1162/dint_a_00255.
- Reus, J.C. (2023) ‘i: gou wei’, *AIMC 2023*. Available at: <https://aimc2023.pubpub.org/pub/hpy32yre>.
- Rich, S. (2021) *Shipwreck Hauntography: Underwater Ruins and the Uncanny*. Amsterdam: Amsterdam University Press. doi: 10.5117/9789463727709.
- Riesch, H. and Potter, C. (2014) ‘Citizen science as seen by scientists: Methodological, epistemological and ethical dimensions’, *Public Understanding of Science*, 23(1), pp. 107–120. doi: 10.1177/0963662513497324.
- Rincon-Gonzalez, L., Warren, J.P., Meller, D.M. and Helms Tillery, S. (2011) ‘Haptic interaction of touch and proprioception: Implications for neuroprosthetics’, *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(5), pp. 490–500. doi: 10.1109/TNSRE.2011.2166808.
- Roach, J. (1996) *Cities of the Dead: Circum-Atlantic Performance*. New York: Columbia University Press.

- Rubinstein, Y. (2020) ‘Uneasy listening: Towards a hauntology of AI-generated music’, *Resonance*, 1(1), pp. 77–93. doi: 10.1525/res.2020.1.1.77.
- Rudin, C., Chen, C., Chen, Z., Huang, H., Semenova, L. and Zhong, C. (2021) *Interpretable Machine Learning: Fundamental Principles and 10 Grand Challenges*. arXiv preprint. doi: 10.48550/arXiv.2103.11251.
- Schacher, J.C. (2010) ‘Seven years of ICST ambisonics tools for Max – a brief report’, *Proceedings of the 2nd International Symposium on Ambisonics and Spherical Acoustics*. doi: 10.5281/zenodo.5184812.
- Schaeffer, P. (1998) *À la recherche d’une musique concrète*. Paris: Éd. du Seuil.
- Schafer, R.M. (1969) *The New Soundscape: A Handbook for the Modern Music Teacher*. Toronto: Berandol Music.
- Schmidhuber, J. (2014) *Deep Learning in Neural Networks: An Overview*. arXiv preprint. doi: 10.48550/arXiv.1404.7828.
- Schnell, N. and Battier, M. (2002) ‘Introducing composed instruments, technical and musicological implications’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.1176460.
- Schroeder, F. and Reuben, F. (2026) ‘Breathing life into AI: Querying AI through tactile imprints using breath and sound’, in Massi, M., Prokúpek, M., Ricci, A. and Ostillio, M.C. (eds.) *Artificial Intelligence in the Cultural and Creative Sectors: Opportunities, Challenges, and Transformations*. London: Routledge, pp. 37–54. doi: 10.4324/9781003509738-5.
- Scolari, C. (2009) ‘The sense of the interface: Applying semiotics to HCI research’, *Semiotica*, 2009(177). doi: 10.1515/semi.2009.067.
- Sconce, J. (2005) *Haunted Media: Electronic Presence from Telegraphy to Television*. 2nd printing. Durham: Duke University Press.
- Scurto, H. and Postel, L. (2023) ‘Soundwalking Deep Latent Spaces’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.11189166.

- Skuggi (2015) *Sorcerer's Creed: The Icelandic Book of Magic Spells*. Reykjavík: Lesstofan. First English edition.
- Shepardson, V., Armitage, J. and Magnusson, T. (2022) 'Notochord: A flexible probabilistic model for embodied MIDI performance', *AIMC 2022*. doi: 10.5281/zenodo.7088404.
- Shepardson, V. and Magnusson, T. (2023) 'The living looper: Rethinking the musical loop as a machine action-perception loop', *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.11189164.
- Smith, A. (1776) *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: W. Strahan and T. Cadell.
- Smith, D.E. and Gordon, A.F. (1999) 'Ghostly matters: Haunting and the sociological imagination', *Contemporary Sociology*, 28(1), p. 120. doi: 10.2307/2653923.
- Smith, H. and Dean, R.T. (2009) *Practice-Led Research, Research-Led Practice in the Creative Arts*. Edinburgh: Edinburgh University Press.
- Stanyek, J. and Piekut, B. (2010) 'Deadness: Technologies of the intermundane', *TDR/The Drama Review*, 54(1), pp. 14–38. doi: 10.1162/dram.2010.54.1.14.
- Sterne, J. (2003) *The Audible Past: Cultural Origins of Sound Reproduction*. Durham: Duke University Press. doi: 10.1215/9780822384250
- Sterne, J. (2019) 'Spectral objects: On the fetish character of music technologies', in Steintrager, J.A. and Chow, R. (eds.) *Sound Objects*. Durham: Duke University Press. doi: 10.1215/9781478002536-005.
- Stiegler, B. (2009) *Technics and Time. 2: Disorientation*. Stanford: Stanford University Press.
- Sturm, B.L.T., Déguernel, K., Huang, R.S. et al. (2024) 'AI music studies: Preparing for the coming flood', *AI Music Creativity (AIMC) Conference*. doi: 10.5281/zenodo.15110181.
- Thagard, P. and Beam, C. (2004) 'Epistemological metaphors and the nature of philosophy', *Metaphilosophy*, 35(4), pp. 504–516. doi: 10.1111/j.1467-9973.2004.00333.x.

- Thompson, R. (2020) ‘Artificial intelligence is powered by ghosts’, *Medium*, 13 November. Available at: <https://sirchutney.medium.com/artificial-intelligence-is-powered-by-ghosts-fe00979914cc>.
- Tomás, E. and Kaltenbrunner, M. (2014) ‘Tangible scores: Shaping the inherent instrument score’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.1178953.
- Torre, G., Andersen, K. and Baldé, F. (2016) ‘The hands: The making of a digital musical instrument’, *Computer Music Journal*, 40(2), pp. 22–34. doi: 10.1162/COMJ_a_00356.
- Tremblay, P.A., Roma, G. and Green, O. (2021) ‘Enabling programmatic data mining as musicking: The Fluid Corpus Manipulation toolkit’, *Computer Music Journal*, 45(2), pp. 9–23. doi: 10.1162/comj_a_00600.
- Truax, B. (2008) ‘Soundscape composition as global music: Electroacoustic music as soundscape’, *Organised Sound*, 13(2), pp. 103–109. doi: 10.1017/S1355771808000149.
- Tsing, A.L. (2021) *The Mushroom at the End of the World: On the Possibility of Life in Capitalist Ruins*. Princeton: Princeton University Press.
- Úlfarsson, H. (2018) ‘The halldorophone: The ongoing innovation of a cello-like drone instrument’, *International Conference on New Interfaces for Musical Expression*. doi: 10.5281/zenodo.1302579.
- Violi, P. (2008) ‘Beyond the body: Towards a full embodied semiosis’, in Frank, R.M., Dirven, R., Ziemke, T. and Bernárdez, E. (eds.) *Sociocultural Situatedness*. Berlin: Mouton de Gruyter. doi: 10.1515/9783110199116.1.53.
- Walshe, J. (2021) *The Text Score Dataset 1.0*. Darmstadt: Darmstädter Ferienkurse.
- Walshe, J. (2023) ‘Ghosts of the hidden layer: Darmstadt, 25.7.18’, in *Choreomata*. London: Chapman and Hall.
- Waters, S. (2007) ‘Performance ecosystems: Ecological approaches to musical interaction’, *Electroacoustic Music Studies Network EMS-07 Proceedings*. Leicester, United Kingdom. Available at: <http://www.ems-network.org/spip.php?article278>.

- West, S.M., Whittaker, M. and Crawford, K. (2019) *Discriminating Systems: Gender, Race and Power in AI*. New York: AI Now Institute.
- Wiener, N. (1961) *Cybernetics, or Control and Communication in the Animal and the Machine*. John Wiley & Sons.
- Wittgenstein, L. (1968) *Philosophical Investigations*. Oxford: Basil Blackwell.
- Xu, F.F., Song, Y., Li, B. et al. (2025) ‘The Agent Company: Benchmarking LLM agents on consequential real world tasks’, arXiv preprint. doi: 10.48550/arXiv.2412.14161.
- Zappi, V. and Tatar, K. (2025) ‘Neural audio instruments: Epistemological and phenomenological perspectives on musical embodiment of deep learning’, *Frontiers in Computer Science*, 7. doi: 10.3389/fcomp.2025.1575168.
- Žižek, S. (2009) *First as Tragedy, Then as Farce*. London: Verso.
- Zotter, F. and Frank, M. (2019) *Ambisonics: A Practical 3D Audio Theory for Recording, Studio Production, Sound Reinforcement, and Virtual Reality*. Cham: Springer International Publishing. doi: 10.1007/978-3-030-17207-7.

10 Appendix

10.1 Publications

Nicola Privato and Thor Magnusson (2025) “The Open Work and AI: Co-creation and Haunting in Text Scores and Inherent Scores”, in *Organised Sound*, 2026:1-11. doi:10.1017/S135577182510085X.

Nicola Privato and Thor Magnusson. (2025) “AI Reads Fluxus: Text Scores, Language Games, and the Art of Being Literal”, in *Proceedings of the International Conference on Technologies for Music Notation and Representation*, TENOR 2025.

https://iil.is/pdf/2025_tenor_privato.pdf

Nicola Privato, Victor Shepardson, Giacomo Lepri and Thor Magnusson (2024) “Stacco: Exploring the Embodied Perception of Latent Representations in Neural Synthesis”, in *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME 2024.

https://iil.is/pdf/2024_nime_Stacco_Exploring_the_Embodied_Perception_of_Latent_Representations_in_Neural_Synthesis.pdf

Nicola Privato and Thor Magnusson (2024) “Querying the Ghost: AI Hauntography in NIME”, in *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME 2024.

https://iil.is/pdf/2024_nime_querying_the_ghost.pdf

Nicola Privato (2024) “Mouja: Experiencing AI through Magnetic Interactions”, in *Proceedings of the 18th International Conference on Tangible, Embedded & Embodied Interaction*, TEI 2024.

<https://doi.org/10.1145/3623509.3635328>

Nicola Privato and Thor Magnusson (2024) “AI Hauntology and the Hauntographic Method”, in *Proceedings of the 12th Conference on Computation, Communication, Aesthetics & X. Fabrica*, xCoAx 2024.

https://iil.is/pdf/2024_xCoAx_AI_Hauntology.pdf

Nicola Privato, Giacomo Lepri, Thor Magnusson and Einar Torfi Einarsson (2024) “Sketching Magnetic Interactions for Neural Synthesis”, in *Proceedings of the International Conference on Technologies for Music Notation and Representation*, TENOR 2024.

https://iil.is/pdf/2024_tenor_sketching_magnetic_interactions.pdf

Nicola Privato, Thor Magnusson, Einar Torfi Einarsson (2023) “Magnetic Interactions as a Somatosensory Interface”, in *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME 2023.

https://iil.is/pdf/2023_nime_privato_et_al_thales.pdf

Nicola Privato and Jack Armitage (2023) “A Context-Sensitive Approach to XAI in Music Performance”, in *Proceeding of the First International Workshop on Explainable AI for the Arts, XAIxArts*, 2023.

https://iil.is/pdf/2023_xaixarts_privato_armitage_context.pdf

Nicola Privato, Thor Magnusson and Einar Torfi Einarsson (2023) “The Magnetic Score: Somatosensory Inscriptions and Relational Design in the Instrument-Score”, in *Proceedings of the International Conference on Technologies for Music Notation and Representation*, TENOR 2023.

https://www.tenorconference.org/proceedings/2023/6TENOR_BOSTON_2023_paper_4288Privato.pdf

Makayla Lewis et al. (2024) “Traveling Arts x HCI Sketchbook: Exploring the Intersection Between Artistic Expression and Human-Computer Interaction”, in *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*.

<https://doi.org/10.1145/3613905.3644069>

Giacomo Lepri, Nicola Privato and Thor Magnusson (2024) “Embodied Sketching for Neural Synthesis”, in *Proceedings of the 19th International Audio Mostly Conference*.
<https://doi.org/10.1145/3678299.3678358>

Jack Armitage, Victor Shepardson, Nicola Privato, Teresa Pelinski, Adan Benito Temprano, Lewis Wolstanholme, Andrea Martelloni, Franco Caspe, Courtney Nicole Reed, Sophie Skach (2023) “Agential Instruments Design Workshop”, in *Proceedings from the 4th International Conference on AI and Music Creativity*.
<https://qmro.qmul.ac.uk/xmlui/handle/123456789/91547>

Giacomo Lepri, Nicola Privato (2024) “Magnetologues”, in *Proceedings of the International Conference on New Interfaces for Musical Expression*.
<https://zenodo.org/records/17553335>

Jack Armitage, Nicola Privato, Victor Shepardson, Celeste Betancur Gutierrez (2023) “Explainable AI in Music Performance: Case Studies from Live Coding and Sound Spatialisation”, in *Proceedings of the Workshop “XAI in Action: Past, Present, and Future Applications”*, part of *Neural Information Processing Systems (NeurIPS) 2023*.
<https://openreview.net/forum?id=se4ojQqjB5>

10.2 Performances

ILL x Spirits in Complexity, with Thor Magnusson and Victor Shepardson. 12 Tónar, Reykjavík. September 2025.

Mouja, InMuse Festival, Vilnius. April 2025.

Mouja+, SWRL Festival, Riga. November 2024.

ILL Diffractive Sonic Dialogues, with Federico Visi, Victor Shepardson, Miguel Crozzoli, Eydís Kvaran, Stefanos Skialivas. Mengi, Reykjavík. October 2024.

Magnetologues, with Giacomo Lepri. NIME 2024, Utrecht. May 2024.

Mouja, xCoAx Conference, Fabrica, Treviso. July 2024.

Mouja, JAM Festival, Venice. June 2024.

Magnetologues, with Giacomo Lepri. TENOR 2024, Zurich. May 2024.

NeoArte: Synthesizer of Arts, with Thor Magnusson, Halla Steinunn Stefánsdóttir, Jack Armitage, Miguel Crozzoli. Nordic House, Reykjavík. March 2024.

Mouja, IRCAM Forum, Paris. March 2024.

Thales, with Jeremy Muller. Guthman Competition, Georgia Institute of Technology, Atlanta. March 2024.

Mouja, TEI Conference, Cork. February 2024.

Mouja, Fabryka Sztuki, Łódź. October 2023.

The Historically Informed AI: Johannes Sperger through Time and Bass, with Darija Andzakovic and Natalia Duarte Jeremias. AIMC, Brighton. September 2023.

Notochord Arcs and Scrambled Signals, with Victor Shepardson. AIMC, Brighton. September 2023.

Mouja, S27, Reykjavík. October 2023.

AI Explainability, Embodiment and Sound in Space, performance lecture at Ars Electronica, Linz. September 2023.

FerroNeural, with Jack Armitage. Instruments, Interfaces, Infrastructures: An Interdisciplinary Conference on Musical Media, Harvard University Department of Music, Cambridge. May 2023.

10.3 Installations and Exhibitions

Feel the Sound, Barbican Centre, London, May to August 2025. With Giacomo Lepri and Halldór Úlfarsson. Presenting a new version of Stacco (V3).

Ars Electronica, Prix Ars Electronica AI in Art Award, Linz, September 2024. With Intelligent Instruments Lab. Exhibiting a new version of Stacco (V2) together with other IIL instruments.

Sónar+D, AI and Music Project Area Exhibition, Barcelona, June 2024. Exhibiting Thales and Stacco as part of the Intelligent Instruments Lab's booth, together with other IIL instruments.

End of Empire at British Textile Biennial, Manchester, October 2023. Large-scale interactive AI installation with Eva Sajovic.

10.4 Workshops

Embodied Sketching with Stacco. Fabryka Sztuki, Łódź. October 2023.

Neural Audio Synthesis and Composition with AI. Jāzeps Vītols Latvian Academy of Music, Riga. September 2024.

Everyday Objects of Musical Improvisation, with Federico Visi, Robert Laidlow, Rebecca Lenton and Theo Nabicht. Humboldt Forum, Berlin. June 2024.

Composing and Performing with AI. Conservatory Pollini of Padua. June 2024.

Sketching Magnetic Interactions with Neural Audio Synthesis. Zurich University of the Arts, April 2024.

Agential Instruments Design, AIMC 2023. With Jack Armitage, Victor Shepardson, Teresa Pelinski, Adan L Benito Temprano, Lewis Wolstanholme, Andrea Martelloni, Franco Santiago Caspe, Courtney N. Reed, Sophie Skach, Rodrigo Diaz, Sean Patrick O'Brien, and Jordie Shier. Brighton. August 2023.

Stacco: Sketching Magnetic Interactions. Fabryka Sztuki, Łódź. October 2023.

Organic Intelligence. With Intelligent Instruments Lab. Hallgrímskirkja, Reykjavík. February 2023.

10.5 Open-source

Thales' specifications, code, and 3D print files are available at

<https://github.com/Intelligent-Instruments-Lab/Thales>.

Stacco's specifications, code, 3D print and laser cut files are available at

<https://github.com/Intelligent-Instruments-Lab/Stacco>.

The RAVE Models trained at the Intelligent Instruments Lab, both by myself and by my colleagues, are available at

<https://huggingface.co/Intelligent-Instruments-Lab/rave-models>.