

ON THE AUTOMATION OF AGENCY IN ALGORITHMIC MUSIC COMPOSITION

A Thesis

Submitted to the Faculty of Graduate Studies and Research

In Partial Fulfillment of the Requirements

For the Degree of

Special Case Doctor of Philosophy

in

Interdisciplinary Studies

University of Regina

By

Jason Roy Cullimore

Regina, Saskatchewan

April, 2021

Copyright © 2021: Jason Roy Cullimore

UNIVERSITY OF REGINA
FACULTY OF GRADUATE STUDIES AND RESEARCH
SUPERVISORY AND EXAMINING COMMITTEE

Jason Roy Cullimore, candidate for the degree of Special Case Doctor of Philosophy in Interdisciplinary Studies, has presented a thesis titled, ***On the Automation of Agency in Algorithmic Music Composition***, in an oral examination held on December 22, 2020. The following committee members have found the thesis acceptable in form and content, and that the candidate demonstrated satisfactory knowledge of the subject material.

External Examiner:	*Dr. Jay Hodgson, University of Western Ontario
Co-Supervisor:	*Dr. Rebecca Caines, Creative Technologies
Co-Supervisor	*Dr. David Gerhard, Department of Computer Science
Committee Member:	*Dr. Craig Gelowitz, Software Systems Engineering
Committee Member:	*Dr. Charity Marsh, Creative Technologies
Committee Member:	*Dr. Helen Pridmore, Department of Music
Chair of Defense:	*Dr. Fanhua Zeng, Faculty of Graduate Studies & Research

*via ZOOM Conferencing

Abstract

This thesis, and the artworks that accompany it, are the products of my Special Case, Interdisciplinary PhD program. The program allowed me to develop artworks where programmable computers are granted agency over music structure, contrasting with my prior composing practice centred on digital audio workstations and linear scores. Here I seek bridges between my established practice of linear composition and those where algorithm-design facilitates interactivity, nonlinearity, and contact between human agents through innovative computer-based interfaces. I establish three types of contribution: one third my contribution to computer science, one third the artworks which I developed, and one third my critical reflection on the process. This reflection centres on my creative methodology, and five main artistic works. These artworks include an original album of songs written from the perspective of hypothetical future, artificially-intelligent machines, three exploratory interactive digital works, and a final major project called “Shards of Memory,” a composition in the form of a smartphone app. This body of work has emerged from my efforts to balance the compositional affordances of programming code with my desire to facilitate the authentic expression of human agents, including myself. I chart my traversal of an interdisciplinary borderland between music composition and computer science, and uncover fertile territories where interdisciplinary artists may challenge musical conventions by using programming languages as compositional media supporting an interplay of multiple agencies. I conclude by arguing that interdisciplinary algorithmic composers are on the vanguard of an evolution of music, and that the chronological evolution of my own doctoral work as an example of how a composer’s practice may be transformed and energized by exerting composerly agency through acts of coding.

Acknowledgments

My supervisors, Dr. Rebecca Caines and Dr. David Gerhard, have provided thoughtful criticism, guidance, and much more, in some instances going “above and beyond” to help me succeed. They encouraged me to find my own voice, provided me with resources to challenge my assumptions, and gave me the freedom to follow my own inspirations.

My committee, Dr. Charity Marsh, Dr. Helen Pridmore, and Dr. Craig Gelowitz, offered contrasting views enriching my interdisciplinary approaches. Dr. Barbara Reul, one of my original committee members, has also consistently shared kindness and wisdom, for which I am most grateful. Dr. Kathleen Irwin and Dr. Christine Ramsay helped to shepherd me through the administrative aspects of my program. Dr. Howard Hamilton helped to procure NSERC funding for the Chief Paskwa’s Pictograph installation, and “Q-Chords” was developed as a project for a class he taught.

I wish to acknowledge all of my friends, with whom I have shared my successes and challenges. Their many kindnesses have helped me immeasurably.

My mother, Susan Cason, has been a supporter of mine throughout my life, and she inspires my compassion and empathy. My brother, Andrew Cullimore, has lent his clear sight, practicality, and good nature. My father, Dr. Roy Cullimore, taught me to look beyond the obvious, and to value the unconventional. I am fortunate to call them my family.

My partner, Max Ferguson, has been a support, an inspiration, and a source of goodness and healing. Our discussions and sharing of our art have been a source of insight and joy. Thank you, Max.

Dedication

My thesis has benefitted from the questioning, affirmation, criticism, and investment of a community of people who have thankfully shared their time and energy with me. So, even though I wrote this text, and designed the projects documented within its pages, I dedicate it to the community of family, friends, and academic colleagues around me whose support and kindness has been so important.

Contents

Abstract	i
Acknowledgments	ii
Dedication	iii
Table of Contents	iv
Chapter 1 Introduction	1
1.1 Overview of This Dissertation	3
Chapter 2 Literature Review	5
2.1 Introduction	5
2.2 What is Agency?	8
2.2.1 Origins and Limitations of Agency	8
2.2.2 Some Psychological Perspectives on Agency	12
2.3 Some Philosophical Perspectives on Agency	15
2.3.1 Structuralist Perspectives	15
2.3.2 Jean Paul Sartre’s Existentialist Perspective	17
2.3.3 Poststructuralist Perspectives	20
2.4 Agency in Music	27
2.4.1 <i>Werktreue</i> and the Authority of the Score	27
2.4.2 Musicking	30
2.4.3 Finding Ourselves In Music Technologies	32
2.5 Relevant Artists	35
2.5.1 Lejaren Hiller and Leonard Isaacson	35
2.5.2 John Cage	35

2.5.3	Iannis Xenakis	37
2.5.4	Roy Ascott	39
2.5.5	Hans Zimmer	40
2.6	Agency Through and Of Machines	44
2.6.1	Software Agents and Frameworks	44
2.6.2	Computer-borne Agency in Human Social Contexts	47
2.6.3	Relating To And Through Machines	52
2.6.4	Augmented Reality	54
2.6.5	The Video Game as a Social Space	56
2.7	Relevant Terminology and Concepts	59
2.7.1	Media and New Media	59
2.7.2	Real-time Composition	64
2.7.3	Algorithmic Music	67
2.7.4	Stochastic Processes	68
2.7.5	Generative, Reactive, and Interactive Music	69
2.7.6	Interactivity	70
2.7.7	Nonlinearity in Music	72
2.7.8	Adaptive Music	74
2.7.9	Comprovisation	76
2.7.10	Soundscapes, Soundwalks, and Audio Walks	79
2.7.11	Mobile and Locative Media	81
2.8	Relevant Works	84
2.8.1	“Voyager:” George Lewis	84
2.8.2	“Ambient 1, Music for Airports:” Brian Eno	86
2.8.3	“Solo” and “Rhythms:” Joel Chadabe	87
2.8.4	“Sonic Meditations:” Pauline Oliveros	88
2.8.5	“Experiments in Musical Intelligence:” David Cope	90
2.8.6	“Aspects of Gaia:” Roy Ascott	93
2.8.7	“Bloom:” Brian Eno and Peter Chilvers	95
2.8.8	“The Cave Of Sounds:” Music Hackspace	97
2.8.9	“Musebots:” Arne Eigenfeldt	98
2.8.10	“Threnoscope:” Thor Magnusson	100
2.8.11	Works of Locative Media	101

Chapter 3	Interdisciplinarity And Methodology	108
3.1	Introduction	108
3.2	Origins of my Practice and Perspective on Music	109
3.3	Research-creation Approaches	111
3.4	Interdisciplinary Approach	113
3.5	Methods	116
3.5.1	Composition In Computer-based Environments	116
3.5.2	Comprovisation	117
3.6	The Computer and the Composer	118
Chapter 4	“We The Artificial:” As Machines May Think	121
4.1	Introduction	121
4.2	Technologies Used to Develop <i>We, The Artificial</i>	123
4.2.1	The Choice to Sing as a Machine	123
4.2.2	Creating a Machine Voice of My Own	124
4.3	Motivations For “We, The Artificial”	126
4.4	Influences on “We, The Artificial”	127
4.4.1	“We, The Artificial” as a Concept Album	127
4.4.2	Influences: “Tommy” and “Computer World”	128
4.4.3	Other Influences	130
4.5	Reflection: Agency in “We, The Artificial”	131
4.5.1	A Depiction of Agency	132
4.5.2	An Enactment of Agency	134
4.6	Conclusion	139
Chapter 5	Q-Chords: Hoarding Agency	141
5.1	Introduction	141
5.2	Overview of Q-Learning	143
5.2.1	A Brief Introduction to Reinforcement Learning	143
5.2.2	Q-Learning	145
5.2.3	Generality of Domains	152
5.3	Towards Q-Chords: Kristopher Reese’s Approach	153
5.4	Q-Chords: A Novel, Q-Learning-based Chord Transition Generator	159
5.4.1	The Domain in Q-Chords	161

5.4.2	Hurrying the Agent	168
5.4.3	Application of Q-Chords within a Video Game	170
5.4.4	Contributions of Q-Chords	171
5.5	Reflection on Q-Chords	173
Chapter 6	Middle Works: A Crisis of Agency	180
6.1	Introduction	180
6.2	List of Works	182
6.3	Moonloops	182
6.3.1	Technology and Function	182
6.3.2	Creation	185
6.3.3	Reflection	186
6.4	Finger Filter	187
6.4.1	Technologies	188
6.4.2	Creation	188
6.4.3	Reflection	189
6.5	<i>Chief Paskwa's Pictograph</i>	191
6.5.1	To Integrate Sound Into a Museum	192
6.5.2	Form of the Installation	193
6.5.3	The Sound Engine And Its Technologies	198
6.5.4	Audio Content	202
6.5.5	Reflection: Problems of Agency	204
6.6	Finding My Voice among Many	206
Chapter 7	"Shards of Memory:" Augmenting Space to Augment	
	Social Bonds	209
7.1	Introduction	209
7.2	Goals	211
7.3	Preparation	212
7.4	Inspirations	212
7.5	Designing "Shards of Memory"	215
7.5.1	Technologies	215
7.5.2	Setting and Gameplay	216
7.5.3	Audio and Visual Content	216

7.5.4	Agents And Agency	222
7.6	Interplays of Agencies in the Algorithmic Music	230
7.6.1	Divisions of Control	230
7.6.2	Conceiving the Algorithmic Score	231
7.7	A Reflection on Mingling Agencies	234
Chapter 8	Conclusion: On Agency and Automation	239
8.1	Agency in Code and Score	240
8.2	An Agentic Mirror	241
8.3	The Agentic Mirror in my Doctoral Works	244
8.4	Harnessing the Agentic Mirror	246
	References	249
Appendix A	“We, The Artificial”: Tracks List and Lyrics	273
Appendix B	Glossary of Relevant Musical Terms	309
Appendix C	Flowcharts Depicting Pathways of Agency	311
Appendix D	“Shards of Memory”: Documentation	335

List of Tables

5.1	Q-functions for a domain of six states.	148
5.2	Example Q-learning reward table.	150
5.3	Example Q-Chords Reward Table	169
C.1	Abbreviations used in Pathways of Agency Diagrams	312
C.2	Corresponding Agencies in Pathways of Agency Flowcharts	313

List of Figures

2.1 Joel Chadabe’s depiction of an interactive composing system.	65
5.1 Illustration of a six-state domain.	146
5.2 Tymoczko’s two-dimensional chord space.	155
5.3 Chord pattern and its trace in Tymoczko’s orbifold.	158
5.4 A Q-Chords domain built for a C-major tonality.	164
5.5 Construction of a domain blending C and E^b -major	166
6.1 Photos of the CPP development space.	200
C.1 Paths of agency A0 and B0: Overviews	314
C.2 Paths of agency A1 and B1: Preproduction	315
C.3 Paths of agency A2 and B2: Cue and memory selection	316
C.4 Pathways of Agency: Cue Analysis	317
C.5 Pathways of Agency: Analysis of a physical space	318
C.6 Pathways of Agency: Overview of theme generation for film	319
C.7 Pathways of Agency: Overview of adaptive theme generation	320
C.8 Pathways of Agency: Determining a new event type in film	321
C.9 Pathways of Agency: Determining a new event type in an adaptive score	322
C.10 Pathways of Agency: Free Comprovisation in film and adaptive score	323
C.11 Pathways of Agency: Constrained Comprovisation in film and adaptive score	324
C.12 Pathways of Agency: Segment duplication in film and adaptive score	325
C.13 Pathways of Agency: Theme duplication in film and adaptive score .	326
C.14 Pathways of Agency: Event evaluation in film	327
C.15 Pathways of Agency: Event evaluation in adaptive scoring	328
C.16 Pathways of Agency: Cue cleanup during film scoring	329
C.17 Pathways of Agency: Cue cleanup during adaptive theme generation .	330

C.18 Pathways of Agency: Feedback from film maker	331
C.19 Pathways of Agency: Mingling agencies in an adaptive score	332
C.20 Pathways of Agency: Delivering a film score after composition	333
C.21 Pathways of Agency: Memory Fulfillment	334

Chapter 1

Introduction

What does it mean when composers, who use digital technologies, entrust a sliver of our agency to a computer by programming it, and so let it act apart from our direct monitoring and control? Where composers choose to do this, what are the consequences? Should we celebrate the technical achievement? Should we fear it? The fact that computers are programmable grants humans agency over them, but, once our agency is so entrusted to the machine, its capacity for autonomous action can detach our agency from our selves, allowing this disembodied agency to act independently of us—or even against us. Questioning of this phenomenon is timely—composers and programmers are already designing systems where portions of human agency are cleaved off to autonomously compose music. The impact of these innovations may shake the world, even as consequential, computer-borne capacities like interactivity, autonomy, and choice suggest that composing with computers might not be a simple extension of past technology-oriented trends, but instead present composers, musicians, and their audiences with something genuinely new. Indeed, the ground beneath the concert halls has already begun to shake, and I feel its rumble keenly.

I am a composer whose pathway into the field is nontraditional. My experience and skills emerge from a combination of independent learning, graduate study of music psychology (at Queen’s University, where I was awarded an M.A. in 1999), and a career that has included commercial and noncommercial projects, leading me to identify as a professional composer and now as an artist working on digital music system design. My embrace of digital technologies, beginning with my first compositions recorded using a Commodore 64 computer in the 1980s, has enabled and energized my

compositional expression. Yet I have worried that the imminent emergence of new automated composing technologies, the descendants of those simple notation programs I used in my first tentative compositions, will ultimately compete with and potentially diminish human composers' opportunities to engage in self-expression. Through my practice of crafting algorithmic musical artworks where I shift the site of my creative expression from the notated, linear score to the dynamic, code-driven music app, I have had the opportunity to explore instances of a rebalancing of human agencies that follows when a work of music is conceived not as a linear document, but as a nonlinear, interactive, algorithmic system. Through this process of research-creation documented in this dissertation, I have constructed algorithmic music systems that trace in their chronology just such a transition. Through their production, I encountered problems associated with the balancing of human- and machine-borne agencies. Each successive work acted to refine my understanding of algorithmic composition as a facilitator of contacts between human agents, and each also reflects my desire to uncover means whereby programming a computer can lead to a magnification of agencies acting musically, as opposed to stifling them.

At a time when automation may be threatening the jobs of workers in many industries, and where artificial intelligences appear in narrative media as harbingers of dystopia or apocalypse¹, I have sought to find myself and others in the medium of algorithmic music, and create algorithmic music that affirms human agency through automation. Through a creative practice that gradually uncovered human agency in automated processes, I have found support for the view that automation of music processes can magnify the creative self, and bring human agents into contact instead of isolating them. In my use of programming environments I have found a medium where algorithms can silence human expression by mediating opaquely, and, in reaction to that insight, I explored a refined approach where digital representations of human agents may flourish in an interplay of agencies. While struggling with the transformation of my own composerly identity, by applying it within the medium of programming code, I found a counterweight to a pessimistic assumption that automation will silence the composer, and an appreciation for the capacity of algorithmic music to magnify humans' self-expressive efforts and promote interpersonal empathy.

1. Examples include the "Terminator" films, the "Alien" films, and the current Star Trek series "Picard."

1.1 Overview of This Dissertation

I will briefly introduce the contents of this dissertation here.

I will provide a detailed review of the literature and creative works that have particularly informed and challenged the development of my doctoral creative projects in Chapter 2. In Chapter 3, I characterize the interdisciplinary underpinnings of my research, and explain the methods which I have used to conduct my creative explorations.

The next four chapters deal with individual creative projects. Chapter 4 introduces “We, The Artificial,” a concept album featuring songs that blend aspects of electronic dance music (EDM), contemporary symphonic film music, and opera. This work is included here as a model of my established compositional practice as it existed before I began my doctoral studies, in a linear form consistent with scored notation. Its composition also acted as my diary, and allowed me to explore the relationship between the conscious mind, technology, and empathy. In Appendix A, I provide a link to audio recording of the album, its lyrics, and brief summaries of the individual songs.

The next three chapters introduce my creative work with coding, and trace the evolution of my understanding of relationships that may exist between automation and the projection of human agency. Chapter 5 describes “Q-Chords”, my first generative music program, which used a Q-Learning framework to automatically compose transitional chord sequences that might be useful as connective material in adaptive music scores such as those found in video games. Chapter 6 details several projects I undertook during my “crisis of agency” phase, a time when I struggled to locate myself in relation to users of the algorithmic music applications I was creating. From this period of questioning, I developed “Shards of Memory”, my final doctoral work, described in Chapter 7. “Shards of Memory” marked a renewal of my compositional practice, where I identify as a composer acting through code to facilitate interactions between human agents, using automation as a means of promoting self-expression and the magnification of interpersonal empathy.

Finally, in Chapter 8, I conclude my dissertation by summarizing its contributions, and reflect upon how my process of exploring the projection of agency through algorithmic music systems might contribute to a dialogue on musical automation between

technologists and sometimes-fearful composers (a group to which I once belonged).

Chapter 2

Literature Review

So you now think of the composer as somebody who plants some seeds,
and then watches them grow.

(Brian Eno, in (Bradshaw 1997))

2.1 Introduction

This chapter surveys relevant literature and practices relating to composition with digital technologies, while exploring related ideas of human and machine agency. I will focus on how human agency has been characterized and projected by artists and technologists who, like myself, conceive of programming code as a compositional medium. I will also explore research concerning the digital computer's capacity to form an artistic medium, a tool, and an autonomous creative partner. I will also introduce key works related to and influencing my own.

The interdisciplinary territories that technologically-adept composers such as myself are now able to explore may challenge established conceptions of the composer, especially in light of the diverse, novel tools and techniques that such interdisciplinary approaches to composition may afford such artists. As such, this review is not narrowly focussed on one primary discipline or line of thought. Rather, I introduce the diverse, relevant perspectives of composers, computer scientists, philosophers, and psychologists, in order to locate my individual practice-based research efforts within an interdisciplinary territory, and to support my argument for the value, and, I believe, the necessity, of conducting such practice-based, interdisciplinary explorations.

The integration of digital technologies into musical activities has shifted the manner in which some composers carry out their activities. As computers' capacities have grown and diversified, and as techniques of software development have been refined, new frontiers for musical expression have appeared. In recent years, for example, "virtual" pop idols have been introduced. The "vocaloid" Hatsune Miku, a virtual construct resembling a fusion of a female Japanese pop star with an anime character, emerges from code and sound libraries that may be purchased and used by any member of the general public (Jorgensen, Vitting-Seerup, and Wallevik 2017, 318-319). Users can, in a common digital audio workstation environment, invoke her voice to sing their own original songs, as if they could puppet a genuine human pop singer. Though Miku has been labelled an "empty vessel" (320), her creator, Crypton Future Media CEO Hiroyuki Itoh, believes she "embodies a vast movement, brought about by the internet, and which blurs the line between creators and users" (Petrarca 2016). This blurring heralds a redistribution of agency, mediated by a novel technology, between the traditionally-powerful music producer and the end user, formerly a "listener" but now potentially a content creator themselves. Thus the virtual Miku's actual "stardom" is an example of how technology can affect human power relationships dramatically.

It is also important to note that both Miku's image and voice have effectively been gendered, despite the fact of their status as technologically-derived entities, rather than human ones. In her 1992 essay "The Gender of Sound," Anne Carson states: "It is in large part according to the sounds people make that we judge them sane or insane, male or female, good, evil, trustworthy, depressive, marriageable, moribund, likely or unlikely to make war on us, little better than animals, inspired by God" (Carson 1994, 119). She goes on to explore representations of (assumed heteronormative) voices and how their differing characters has been used to support oppressive patriarchal practices. While this topic is outside the scope of this thesis, it can provoke a questioning of how sound, and apparatuses used to produce it, might affect views of music-related technologies by projecting them through a filter of gender. Numerous researchers have identified instances where musical technologies have been referred to in gendered terms. Paul Theberge, for example, observes digital technology being cast in feminine terms to appeal to male purchasers (Theberge 1997, 123-124), and Jill Halstead and Randi Rolvsjord examine how the gendering

of electric guitars marketed to women (through modifying a typical guitar’s size or colour) could affirm negative, passive views of femininity (Halstead and Rolvsjord 2017, 15). This dissertation is not directly concerned with contributing to a feminist critique of algorithmic music composition. Nevertheless, when human agencies come into contact or conflict with one another, a power dynamic can be introduced that may relate to disparities of agency identified by feminist scholars and artists. Hence, while I do not deal with feminist perspectives on power in great depth here, I am concerned with power balances between agents in contexts of music-oriented social interplay, and will introduce some core ideas from feminist research in my discussion.

Software emulating the action of human composers has existed for decades, and its refinement might prompt us to ask if we might soon place the name of an algorithm alongside venerable composers such as Ludwig van Beethoven or John Williams. Such a development seems imminent: in 2019, for example, an automated composition system known as “AIVA”¹, went into an open beta test. Marketed as an aid for composers, it has also distinguished itself as the first software program registered as a composer with a royalty-collecting Performing Rights Organization (here, the French PRO SACEM) (Li 2018, 355). Established technologies such as physical distribution media (e.g. CDs and cassette tapes) and time-worn business models (e.g. the record label acting as a gatekeeper, the radio as a promotional venue), are being reimagined, and sometimes supplanted outright, by alternatives like subscription-based music streaming and social-media-driven “do-it-yourself” career paths supported by digital environments (Thomson 2013, 515). The transformative relocation of music-making activities into the digital domain challenges the composer to adapt, as disruptions to established technological and cultural systems supporting music-making impact upon how composers, musicians, and audiences make and share music. On this musical ground, where human agency is automated and projected through algorithms, I situate my dissertation and its associated creative work. I believe that this ground, uncovered by engineers, scientists, and new media artists, presents challenges and opportunities for transformative evolutions of human agency that may affect how we engage with, and even define, music. As such I begin this Literature Review with an exploration of the concept of agency, through the lenses of artists and composers, scientists like psychologists, and philosophers.

1. Official website for AIVA: <https://aiva.ai/engine>

2.2 What is Agency?

The idea of “agency” is found in many disciplines, although it may take on slightly different connotations in each. In this section I will present several disciplinary perspectives regarding the idea of agency, beginning with a general definition, and then introducing disciplinary definitions more familiar to composers and computer scientists.

Contemporary definitions of agents and agency can be reduced to essential qualities. An “agent” may be understood as “a being with the capacity to act,” while the term “agency” refers to “the exercise or manifestation of this capacity” (Schlosser 2019). An agent may also be a vessel for another’s agency, a distinct person or thing through which power is exerted or an end is achieved (“Rationality” 2015).

Some key facets of agency are apparent in these characterizations. There is reference to power being projected. There is a vessel containing or transmitting that power. There is autonomy, in that the vessel holding another’s agency may act independently of them. In this sense, agency is something we project outward from ourselves. On the other hand, agency can also be considered an intrinsic ability: “the human capacity to act” (Ahearn 1999, 12). All of these capacities are consistent with my own conception of agency.

2.2.1 Origins and Limitations of Agency

Laura M. Ahearn provides a useful perceptive on the origins of the study of agency. As a subject of scholarly study, agency gained currency in the late 1970s, at a time when activism sought to challenge power structures in support of race and gender equality (12). Feminist theorists have explored how people’s actions might affect social and political structures, and vice versa, with examples including Banu Ozkazanc-Pan’s analysis of agency in the “#MeToo” movement (**Ozkazanc2019a**), and Susan Clegg’s article exploring agency in the context where poststructuralism encounters feminism (Clegg 2006). Clegg introduces her text by arguing that agency, whether individual or collective in form, is “is at the heart of the feminist, and indeed, all radical political projects” (309). For her part, Ahearn cautions that the term has, through much of its history, been of nebulous definition, which is a concern since use of the term may have implications for our understanding of “personhood, causality,

action, and intention” (Ahearn 2001, 112). Thus, when discussing agency, we are likely discussing something precious.

Ahearn provides a provisional definition that she acknowledges is incomplete: “agency refers to the socioculturally-mediated capacity to act” (112). This provisional definition, Ahearn explains, is merely a starting point, and she points out that it fails to adequately account for questions that locate agency outside of an individual within a society (a group of people, or an object, for example). Of particular interest in this dissertation is the question as to the nature of agency in machines—whether agency might fall entirely in the purview of the human individual, whether it might be shared with or contained within an autonomously-functioning algorithm, or whether agency might be super-individual, projecting the action of multiple individuals, so that they might impact upon one another.

Arguments against the existence of human agency have also been raised. Ahearn gives the example of Michel Foucault, whom she says can be read both as supporting or denying human agency entirely (116). Ahearn notes that in his book “History of Sexuality, Volume I,” Foucault allows that power evokes resistance, an opposing force, but does not clearly explain how either power or resistance might be enforced or personified within or between individual humans (Ahearn 2001, 116; Foucault 1978, 93,95). The idea of power relations is germane to this thesis, since, arguably, coding may involve a projection of power through code, creating a system of power relations that may benefit or do harm to software users. I believe that such power relations can be viewed through the lens of governance, since I view a program as a system of constraints devised by a programmer to channel its users’ actions or thought towards some end (stated or unstated).

To Foucault, to exert power is to “guide the possibility of conduct,” a synonym for leading (Foucault 1982, 789). There are three forms of conduct: the state of conducting another person, the state of being conducted by another person, and the conducting of oneself (Foucault 1984, 257-58). I introduce this concept here because the environment of a computer program—its interface and what it allows one to experience and do when interacting with it—can be viewed as the programmer enforcing the conduct of users through code design, limiting users’ options in a manner that makes the user subject to the programmer’s will. In Foucault’s view, governing someone involves structuring their field of freedom, as opposed to “reigning,” “ruling,”

or “commanding” them (Foucault 1984, 161). Power is not equivalent to total control, which negates freedom. Rather, the exercise of power is actually an attempt to conduct the conduct of another human agent or agents (Foucault 1982, 790). Human government, then, evokes a tension between conducting the self, conducting others, and being conducted by others, setting up a conflict between autonomy and the choice to submit to external command. In the conflict between agents that ensues, the scope of an individual’s freedom is identified, defined and structured on the basis of the choice to accept or refuse to be conducted by whatever particular mechanism is acting upon them (Lorenzini 2016, 10).

I believe that the exercise of power manifesting through designing an algorithm can essentially be an act of attempting to conduct the behaviour of others (e.g. the users of the program that interact with the algorithm). This capacity to conduct through the vessel of code is an element of the nascent power of the programmer: the shaping of constraints upon the user’s conduct within the bounds of the program’s possible actions. “Counter-conducts”—acts of resistance to the prevailing conduct under which one may be constrained—is not a rejection of control, but an expression of the human agents’ desire to be governed under a different form of conduct (11). When I examine interactive, algorithmic music systems, I am most interested in the degree of agency that each party in the realization of such a work possesses in the effort to realize the work. Foucault’s idea of conduct is useful when analyzing or categorizing the quality of an algorithmic musical work because it relates to the flow of action a user takes when engaging with a program—to be taken down a branching story path that offers some options but not others; to make an avatar that has one characteristic but not another; to be allowed to trigger a major harmony but not a minor one. Choice occurs within such a framework of conduct, where particular systems of conduct enforced by the programmer may be embraced or resisted. As I progressed through the works presented in this dissertation, I was (unknowingly) seeking means of balancing the three forms of conduct, which I framed as balancing an interplay of distinct agencies (e.g. those of a user, a content contributor, or myself). Agency might be seen as a prerequisite for freely projecting power or resistance to power, but conduct relates to agency’s expression outside of the individual, and its effect(s) in the world around us. As such, the medium in which the composer works (e.g. programmatic versus notational) may be subject to the same agency acting through it, but, because

different mediums offer different systems for enforcing conduct, they diverge in how the agent may act through them. I will return to Foucault's ideas in the upcoming section dealing with poststructuralism (Section 2.3.3).

Another debate regarding the existence of agency has originated from its association with the concept of free will. In the context of philosophy the term "free will" can refer to the the human capacity to act, or not to act, as we so choose, apart from compulsion or restraint applied by forces external to the individual (Shogren et al. 2015, 256). To have agency implies that when one acts, one could have chosen to do otherwise (Feldman 2017, 2). Most people experience their existence as "conscious agents," assuming they can exert conscious control over their actions, a force emanating from the individual self (Mascolo and Kallio 2019, 437-438). The nonexistence of free will might threaten a conception of agency that depends on it. Even the appearance of constraints on free will may appear to limit agency, be they internal (e.g. one's intelligence, genes, or disability) or external (e.g. other people's agency, behavioural norms, natural forces, or even God)(Feldman 2017, 2).

Agency is often thought to require a capacity for choice (5). Determinism, a philosophical doctrine holding that all events (including human actions) are determined by preexisting causes (Shogren et al. 2015, 256), seems to negate both free will and agency by denying humans the capacity to choose. One might expect determinism to cause fear among powerful institutions such as government or religion, for if a population believes their actions to be preordained, they may think that their behaviour is inevitable, and that they are not truly responsible for it. This in turn might make it easier for them to engage in behaviours that are not socially-sanctioned or dangerous to the institution. Spreading the belief in free will, and the existence of human agency, might be viewed as a beneficial "social tool" for institutions to promote. By so doing, they ensure individuals' belief in their accountability for their own actions, which may be "essential to legal, moral, and political judgements" (Feldman 2017, 3). Here, agency paradoxically becomes an enabler of social control, by forcing the individual to more deeply consider the consequences of his or her actions.

Foucault describes this effect in his discussion of the "panopticon," an Eighteenth-century innovation in prison design proposed by Jeremy Bentham. Bentham's panopticon was a prison in which inmates are allowed relative freedom of action within the building, but only with the knowledge that they may be observed by the guard at

any given time (Foucault 1995, 199-201). To Foucault, the government's aim to see and thus know all that people do is the foundation of the modern state's aspiration to control (223). The internet itself has been identified as a panopticon, where the threat of constant surveillance has reportedly prompted many users to exercise power over themselves through choosing to avoid certain activities (e.g. purchasing particular items, or visiting certain websites), based upon the assumption that their activities may at any time be surveilled by authorities or people with whom they share on social media (Manokha 2018, 232). Since the internet is sustained by digital computers running algorithms, and these algorithms have the capability of conducting the conduct of others through threat of being watched, the evocation of Foucault's panopticon can be a feature of a program, causing users to conduct themselves differently in accordance with their fear of being seen.

Programmers such as myself who wish to engage other individuals' agencies through code, then, should recognize that their position is privileged, giving them power within the confines of the program to affect others' conduct in ways that those others cannot replicate. Free will, even if it exists, may thus be constrained in software users simply because they recognize that their activities may be observable, interpretable, or recordable by the code. In this sense, there is a power imbalance inherent in creation and consumption of software that may mitigate the gains in self-expression associated with empowering features of software such as interactivity or nonlinearity. And, where the user recognizes this, free will becomes less free, even where it does actually exist. In this sense, the coder's privilege may resemble the composer's, in that there is a boundary that separates creator from consumer. It is defined by the ability to specify notes on the part of the composer, or functions on the part of the coders. The authority of the notated score may be reinvented in the authority of the algorithm, an authority from which the musician, audience, or user might be barred, no matter how freely they may wish to act.

2.2.2 Some Psychological Perspectives on Agency

In the sciences, other approaches to understanding agency have emerged. On one extreme is a repudiation of agency advanced by behaviourist B.F. Skinner. Skinner's studies of animal behaviour relied upon objective, replicable measurement of animal

responses to stimuli (Morf 1998, 35). To Skinner, humans were functionally no different from primates, or even inanimate objects such as rocks, in that all are subject to a common, pervasive determinism that he believed was powering all processes of change (35-36). That human behaviour could be spontaneous, arising from some internal spark that reveals free will or a soul, was out of the question to Skinner. He articulated this point by recalling the idea of the spontaneous generation of maggots and microbes on meat, an idea that was experimentally dispelled by Louis Pasteur in a mirroring of behavioural science's more recent dispelling of the fantasy of free will (35). This line of thought has been described as an "exhaustive materialism," one that treats free will and the spark of an internalized, human agency as no more than illusion (36).

Other psychological perspectives are not so deterministic, allowing room for agency to exist. In his article "Towards a Psychology of Human Agency", Albert Bandura defines agency through the lens of social cognitive theory (Bandura 2006). "To be an agent," he writes, "is to influence intentionally one's functioning and life circumstances." (164) He ascribes four core properties to human agency (164-165):

- Intentionality, including devising plans for action and strategies for achieving them;
- Forethought, a temporal extension of agency into the future that provides "direction, coherence, and meaning to one's life;"
- Self-reactiveness, which supports planning by permitting the individual to "construct appropriate courses of action... and regulate their execution;" and
- Self-reflectiveness, which is a capacity to examine one's own efficacy and adjust behavioural strategies in response, should such adjustments be deemed necessary.

Bandura argues that humans are not wholly autonomous agents, nor are we utterly reactive to environmentally-based influences. Rather, he argues, our functioning is a product of a reciprocal interplay of what he calls "intrapersonal, behavioural, and environmental determinants" (165). In this view, humans not only cause events, but are also subject to them, and at points of contact between experience and intention we find the forces powering our behaviour.

The concepts of intentionality and autonomy are both closely related to agency. Bandura writes, “an intention is a representation of a future course of action to be performed. It is not simply an expectation or prediction of future actions but a proactive commitment to bringing them about.” (Bandura 2001, 6). Intent is located in contexts where the agent understands the connection between his or her action and the outcome, and where he or she acted deliberately to shape it (Feldman 2017, 10).

“Autonomy”, according to Feldman, refers to “the self as maintaining a separate and independent self from other agents, without addressing the many other types of constraints to free will, either external or internal” (10). The idea of autonomy has been considered a defining feature of a living organism (Shogren et al. 2015, 256). It also appears prominently in discourses of computer science. A software agent, for example, may be categorized thus on the basis that it is capable of autonomous, self-directed behaviour (Mostafa et al. 2017, 9). A more thorough discussion of software agents appears in Section 2.6.

The preceding discussion touches on several core ideas relating to agency, from multiple disciplinary perspectives. The conceptualization of agency as presented in this dissertation, however, needs to address questions such as:

- Can a piece of software contain or project an internalized agency?
- What happens to a composer’s agency when they create through the medium of program code, and how might this relate to acts of writing a score, whether that score is linear, incorporates chance elements, or elevates a performer’s agency over notational expression through allowing them to improvise?
- How might digital technology support the development of environments where human agencies can interact and generate meaning?
- How will the ongoing evolution of our understanding and practice of music be affected by such technological developments?

In the remaining sections of this literature review, I will approach these questions through a deeper examination of agency as it may be understood in the domains of philosophy, music, and computer science. I will also introduce representative art works and music that embody or challenge these ideas, and that have impacted upon my own artistic practice. I will begin with a brief exploration of agency in philosophy.

2.3 Some Philosophical Perspectives on Agency

My purpose in this section is not to present a history of philosophy, but to indicate how various philosophical perspectives can expand my characterization of agency.

Addressing this topic is beneficial because the creative work I present in this dissertation may challenge some understandings of human agency by applying the concept to algorithms. If, for example, we argue that agency is constrained by external systems of social relations as opposed to originating from some human spark of intent, what does that mean for an algorithm’s choices? Do we ascribe it an agency that it carries out as if it were a social being like another person? Or, is an algorithm’s agency the imprint of its designer, mirroring in the execution of her code her own social status? Such questions can be identified through exploring contrasting conceptions of agency by philosophers. I will briefly discuss human agency as seen through three lenses: those of structuralism, poststructuralism, and Sartre’s existentialism. My accounts of these ideas are not exhaustive histories, but rather are intended to ground my perspective on agency as it is projected towards, through, or by machines controlled by algorithms.

2.3.1 Structuralist Perspectives

In simple terms, a “structure” may be considered a “relationship between variables” (Dowding 2008, 24). Structures can often take multiple forms, even when their components are in common. An example might be found in architecture, where a bungalow might be made of bricks that otherwise could have formed a two-story house (24). In music, the same set of pitches could be used to form a particular lullaby, symphony, or transcription of a birdsong, and the score forms a structure in which the pitch-variables are arranged in different ways that constitute distinct compositions.

What if we were to substitute bricks or pitches with people? The structure formed from groups of people would be a web of social relationships. And, like the bricks forming houses, the same people might form very different social structures in different iterations—a hierarchy, or an egalitarian society, for example (24). The form of a social structure, enforcing as it does a distinct set of roles for its constituent individuals, can be understood as a system of power relationships. Can an algorithm

play a meaningful role in such a human system? The question can be approached from a structuralist perspective.

A structuralist would describe a power structure by examining the relationships between the people who act within it (Dowding 2008, 24). Claude Levi-Strauss, a foundational figure in the development of structuralism (Ortner 1984, 135; Ruitenberg 2018, 694), sought to clarify the bewildering chaos of culture, as seen through the lens of anthropology, by adapting the methods of structural linguist Ferdinand de Saussure (Levi-Strauss 1963, 20-21). Saussure argued that meaning in language was relational, deriving from differences and similarities between terms, rather than from inherent meanings embedded within each term (De Saussure 1965, 114-15). To the structuralist, it would be an error to assume that the meaning of a term comes directly from the object to which that term refers (Harcourt 2007, 4). Rather, meaning falls from the relationship between the constituents of language, which Saussure labeled “signs” (De Saussure 1965, 15). Signs are combinations of signifier (e.g. a marking or assembly of phonetic sounds), and signified, which corresponds to the concept associated with the signifier (66-67). The relationships between signifier and signified are arbitrary, and to ascribe meaning to a concept is to understand it in terms of related signs (67). Language becomes a system where each element depends upon the coexistence of all others to attain meaning, a process Saussure argued was carried out in the unconscious mind (72).

Why is structuralism introduced here? It is because this work is also about certain types of structures—programmatically structures that govern the functioning of algorithms, and social structures which organize the system of social relationships that form around the execution of an algorithm. This is not to say that structuralism provides a template for creativity. In his article “Structuralism and Myth,” Levi-Strauss argues that:

...too many contemporary works, not only in the literary field but also in those of painting and music, have suffered through the naive empiricism of their creators. Because the social sciences have revealed formal structures behind works of art, there has been a rush to create works of art on the basis of formal structures. But it is not at all certain that these artificially arranged conscious structures are of the same order as those which can be discovered, retrospectively, as having been at work in the

creator’s mind, and most often without any conscious awareness on his part. (Claude Levi-Strauss, in (Levi-Strauss 1981, 68-69))

Levi-Strauss is arguing that structural analysis is best not considered to be generative of art or music, and thus constraining composers to a particular framework. Indeed, in a subsequent passage he states that “instead of composing new music with the help of computers, it would be more relevant to use computers to try to understand the nature of existing music” (69). This is, as might be expected, a view I find short-sighted, in that I believe it assumes that computer-generated score does not form a creative act, which, for the coder, it is. Nevertheless, he rightly argues that composers might find in structural analysis a “new awareness” of the foundations of music—the understanding of which can spark innovation (69).

This is an important idea, illuminated in his analysis of myth: “Myths tell us nothing instructive about the order of the world, the nature of reality, or the origin and destiny of mankind... On the other hand, they teach us a great deal about the societies from which they originate... and most importantly, they make it possible to discover certain operational modes of the human mind” (66). These “operational modes of the human mind,” to Levi-Strauss, are constant in time and across geography, and so understanding them reveals truths that link all humans (66-67). My own practice of composing through coding is predicated on the idea that machine- or software-borne algorithms can be regarded as social actors coexisting among humans making music. While I am not a strict structuralist, I appreciate how, in uncovering common principles governing human and machine agency, we arguably begin to join them in a structure like a society, an idea I will return to in Chapter 7.

2.3.2 Jean Paul Sartre’s Existentialist Perspective

We now turn to Existentialism, introducing a few ideas of the French existentialist Jean-Paul Sartre, in whom structuralism found a contrasting view where human agency emerges from its obscure position in the interrelations of culture, and returns boldly to its seat at the centre of the human mind.

In Sartre’s view, the structuralist signifier was not a term in an arbitrary association with a signified concept. Rather, Sartre asserted that signification is a process he enacted himself—that by naming a thing, he, Sartre, became its signifier. In this

view, the human actor is an agent who is “nothing other than what he makes of himself” (Sartre 2007, 10). Meaning here originates in the subject’s personal experience and behaviour, beginning with the essential Cartesian axiom “I think therefore I am,” and extending through each subsequent action that a human agent carries out (Harcourt 2007, 9-10). Sartre states: “Man is not only that which he conceives himself to be, but that which he wills himself to be, and since he conceives of himself only after he exists, just as he wills himself to be after being thrown into existence, man is nothing other than what he makes of himself.” (Sartre 2007, 22). Perhaps we all yet have the power to make our own realities.

In Sartre and Levi-Strauss’ opposed perspectives we find diverging characterizations of the limits of human agency. For the structuralists, the relations of signs operate in individuals’ unconscious, out of their conscious control (Harcourt 2007, 5). Since the laws governing language are beyond the control of individual speakers, they may be studied as objective phenomena to explain human behaviour (5). Sartre, on the other hand, not only asserts the existence of free will; he has been characterized as “the most extreme exponent of the position that human beings are absolutely free to choose between options” (Morf 1998, 31). In this sense, he is located at the opposite extreme from Levi-Strauss and of B.F. Skinner, for whom the mind was an inconsequential black box bereft of a meaningful agency.

To elaborate on Sartre’s conception of agency, we may consider his two modes of being, as illustrated by the example of a hiker and a rock (the following scenario is adapted from a discussion of Sartre’s ideas by Martin E. Morf (32-34)). Sartre describes the first mode of being, “being-in-itself,” as applying to objects unconscious and incapable of being anything other than what they are (Sartre 1993, lxv). Such a being can be possessed by an inanimate object, such as a large rock. The rock can be considered “neutral,” utterly incapable of self-determined movement. This quality that defines beings-in-themselves, for Sartre distinguishes such beings-in-themselves by their inability to direct themselves to change—they simply are what they are (lxv). Yet beings-in-themselves, despite their inertia, can still create conditions around themselves about which a human agent can do nothing, such as blocking a hiker’s path. To have the capacity to block the hiker’s path in a way that defies the hiker’s agency (assuming that she cannot shift the rock) means that the rock possesses the quality of “facticity.” To have facticity means that the object creates

conditions about which an agent can do nothing. Another element of the scene that would possess facticity is the personal history of the hiker, which is locked in time and thus cannot be changed by her. That neither the rock nor the hiker's past can be changed might seem to imply that the hiker lacks agency over either of them. But Sartre does not see facticity as a nullification of agency. Instead, he addresses human agency by introducing a second mode of being that he terms "being-for-itself."

While being-in-itself and being-for-itself are concepts in opposition to one another (Sartre 1993, lxiii), they also exist simultaneously in the human, for she is both a facticity and a "transcendence" (56). It is a facticity, because the human will is limited, for instance, in choosing its own identity—otherwise one might say that she is richer or poorer than she actually is, and by so doing make it so (83). Yet the being-for-itself is also conscious, and so can change a being, a process of nihilation that affects a particular being-in-itself (618). So, the capacity to transcend the inertia of a being-in-itself cannot magically grant the hiker total agency over the state of the obstructive rock, nor over her past. It does, however, grant the hiker a freedom to choose the meaning of such facticities. The rock, though immobile, might be viewed by the hiker as a restrictive obstacle that inhibits all desired movement, as a challenge to be climbed, or as merely an invitation to find another, more agreeable path. The history of the hiker is likewise utterly determined, but the hiker is nevertheless free to view that facticity in her own terms, a process of transcendence that allows the conscious being to surpass the facticity of the rock that grants the conscious entity its freedom (484-85).

This concept of agency might be applied to an algorithm, raising the question as to whether an algorithm is essentially a being-in-itself, or whether it can meaningfully contain the agency of a being-for-itself. I will approach that question in Section 2.6. As for Sartre, his opposition to a core principle of structuralism has been framed as an instance of a broader dichotomy between humanist psychology and psychology positioned as a natural science (Morf 1998, 36-37). I will return to this topic in Chapter 3, which outlines my key methods and interdisciplinary perspective. For now, I will continue by bringing the discussion of agency closer to the present day, introducing poststructuralism.

2.3.3 Poststructuralist Perspectives

Poststructuralism is not strictly an opposition to structuralism (Ruitenbergh 2018, 694). In fact, it builds on some structuralist notions, such as the idea that meaning derives from relations of signs operating in the subconscious (Harcourt 2007, 17). But where Levi-Strauss sought universal principles underlying structures of meaning, poststructuralists looked at ambiguities in systems of meaning, and found meaning within them (17). As Ruitenbergh says, “one of poststructuralism’s significant departures from structuralism is that it shifts its focus on what language means to what language does” (Ruitenbergh 2018, 694).

Several poststructuralist philosophers have relevant perspectives on agency. The full study of these perspectives is not a concern of this dissertation. However, I will still introduce two ideas, outlined in very brief terms, as they have influenced my own perspective on agency, followed by a more detailed account of Roland Barthes’ essay, “The Death Of The Author.”

Michel Foucault: Language as a Prison

Michel Foucault depicts the human being as a “caged animal” who is caught in the prison of her language, where language is the only means through which a human can access reality (Jeongwon and Song 2002, 265). Foucault is identifying language as a filter constraining us from knowing the absolute truth, which we can never touch directly. I see this idea in terms of a denial of an artwork’s capacity to transmit meaning from artist to audience—like the mp3 audio format, language is “lossy,” conveying meaning through a compression of information that discards unnoticed elements, as well as those which the unconscious functioning of the mind discards due to the limits of cognitive processes. A composer’s agency over a score already forms a splintering of expression, as notated scores are reductive, leaving out precise and continuous variations in pitch, dynamics, tempo and articulation that are normally reconstructed in the form of the musician’s interpretive performance. And still this performance is filtered by the listener, first in the eardrum, then in the cochlea, and finally in the many neural processes occurring throughout the brain. I can no more know Beethoven through reading or hearing his music, as I can be assured that the painting on the wall currently hanging above me is of a “moose.”

To claim otherwise is arguably an act of faith.

Deleuze and Guattari: The Rhizome

Gilles Deleuze and Felix Guattari introduced rhizomic thinking as a reconsideration of how we think (Sutton and Martin-Jones 2008, 4) in their book “A Thousand Plateaus: Capitalism and Schizophrenia” (Deleuze and Guattari 1987). We may at times find ourselves reducing things to oppositional binaries, such as the composer and the listener, or the performance and the score, with the result being a hierarchical organization of ideas. Deleuze and Guattari liken this model of thinking to a tree, what they call the “image of the world,” a classical book, the hierarchy of linguistics (5). A tree has an origin (the seed), a collection of distinct features with proscribed relationships (e.g. roots that anchor and nourish, a trunk that gives height, and leaves that form the sugars that power metabolism), and a distinct separation from those surrounding elements that are not the tree (e.g. the tree is not air, ground, nor forest). The arborescent model is applied to many human innovations, whether technological or social. For instance, a corporate hierarchy is a tree, as is the Western tonal system of music, for both are hierarchies (Spyrou 2019, 37-39), one headed by a CEO, the other by the tonic pitch. The rhizome is found in an alternate perspective. “A rhizome as subterranean stem is absolutely different from roots and radicles,” they state, referring to the metaphoric tree they have introduced (Deleuze and Guattari 1987, 6). Like the fungal network that shares the name, the rhizome is all-encompassing: “any point of a rhizome can be connected to anything other, and must be. This is very different from the tree or root, which plots a point, fixes an order” (7). The rhizome is a pack of rats, or the burrows dug by rabbits (6). It is in music, as when Glenn Gould accelerates the tempo of his performance, “transforming musical points into lines.... making the whole piece proliferate” (8).

While describing the Western tonal system of music as an embodiment of arborescent thought, Deleuze and Guattari see music as a strongly rhizomic art:

By placing all its components in continuous variation, music itself becomes a superlinear system, a rhizome instead of a tree, and enters the service of a virtual cosmic continuum of which even holes, silences, ruptures, and breaks are a part.... the important thing is certainly not to

establish a pseudobreak between the tonal system and atonal music; the latter, on the contrary, in breaking away from the tonal system, only carried temperament to its ultimate conclusion.... The essential thing is almost the opposite movement: the ferment in the tonal system itself that dissolved temperament and widened chromaticism while preserving a relative tonality, which reinvented new modalities, brought a new amalgamation of major and minor, and in each instance conquered realms of continuous variation for this variable or that.... Music is not alone in being art as cosmos and in drawing the virtual lines of an infinite variation (Deleuze and Guattari 1987, 96-97)

Deleuze and Guattari sought to replace the tonal system with a “generalized chromaticism” where every component sound, such as “durations, intensities, timbres, [and] attacks,” is placed in “continuous variation” (Campbell 2013, 701; Spyrou 2019, 39). An example of a work conceived as a rhizome is the English-language opera “Faustus, The Last Night” (2006) by French composer Pascal Dusapin. The composer crafted the music and libretto simultaneously, assembling the text from scattered sources (e.g. memories of books, clippings from newspapers, offhand remarks directed at others), not following a grand design or plan (perhaps concerned that a narrative structure would impose an undesirable framework), and letting the creative process proceed without a “central idea” (Spyrou 2019, 39-40). The composition could be said to have “grown” more than “developed,” a reflection, perhaps, of its composer’s choice to deemphasize preplanning and avoid applying a formalized structure to govern the process of its creation (42-43). Of course, the opera did require planning in order to mount it successfully, and it would be difficult to argue that the individuals who mounted it could cast away their experience with music as it formed. Indeed, music can be consistent with both a tree and a rhizome. The hierarchy of notation can be found in a web of relations between pitches, degrees of metric emphasis, and even the stratification of the social environment of the classical concert or modern music festival. Yet where such conventions are ignored or subverted, rhizomic musical qualities may be revealed, in the blurring of roles and decentralization of control of the jam or group improvisation, or in the graphic score that facilitates musical performances through ambiguous symbolic prompts rather than more unambiguous

markings governing pitch and instrumentation² .

The computer, in its highly-ordered programming languages and hierarchically-organized component systems, seems to be an embodiment of the tree. Yet the behaviour of the computer need not be so constrained. It can behave unpredictably, whether through engaging in processes of randomization or acting in accordance with a behavioural model that its users cannot predict. Such behaviour might seem to be unhinged from its coder, appearing to be an agency that, in its detachment from perceived human control, might appear alive. I feel that this form of emergent behaviour is a contribution of the rhizome to algorithmic music. The rhizome provides an alternative to a system that is ordered. Its appearance might resemble the fall of a classist society, where one has learned to act in accordance with their perceived social position—their place in a hierarchy. The rhizome dissolves the hierarchy, in turn approaching some state of anarchy. An anarchic state can be a place of anxiety, where the chaos of unstructured action poses a threat. Yet it is also a state where one can be free if one chooses to fight to be free, for in the traversal of a rhizome, the freedom to go where one wants, and do what one chooses, is magnified by the dissolution of the constraints of the hierarchy. For the computer musician, the rhizome offers a break from a socially-constructed roles of composer, musician, audience, and so forth. In its place, agents encounter one another and blend in each others' eyes. It is ironic that the computer, which typically depends upon the most rigid, unambiguous specification of coded instructions, can help the composer form works that sweep aside the established social and musicological traditions that govern both music and its makers. Computer-sustained rhizomes, such as the world wide web or open-ended generative score, are possible because, paradoxically, the computer can be deliberately and intricately programmed to serve chaos, to make rhizomic connections between all things and each one of us. The fostering of an algorithmically-based musical rhizome is a celebration of music as a connective force. Such a rhizomic music ranges widely and unpredictably, and seeks to draw in any humans in the range of its

2. It is also important to acknowledge that the rhizome has been a concept that found synergy with new media art and artists. The community surrounding and supporting the website Rhizome.org, for example, was originally formed by Mark Tribe in 1996 as an online project centred around the development and preservation of new media and net art. The organization emphasized grassroots action outside the institutional domains of established galleries, museums, and publications, which could arguably be seen as a reaction against the hierarchical structures which Deleuze and Guattari critiqued and challenged through their decentralized, networked rhizome.

grasp. Such music echoes the drum circle, the communal improvisation, and ecstasy of freely making music and dancing with one another. In this sense, a rhizomic computer music does not take us into new territories, but re-establishes a more ancient music: disorganized, spontaneous, and direct.

Roland Barthes and Birth of the Reader

Roland Barthes' essay "The Death of the Author" (Barthes 1977) challenges the notion that a literary work contains and transmits the essence of its creator's original thoughts and history. Barthes suggests that, in his time, the evaluation of literary works was distracted with questions pertaining to the nature of those works' authors, and those authors' identities had come to be viewed as the secret meaning underlying their work (143). Under this paradigm, the purpose of reading had become the location of the author, and, once the author is identified, the work is considered effectively "explained" (147).

Barthes argued against such a framing of literary interpretation, noting that attempts to discern the author at the root of his or her writing actually limits it, casting it as merely an artifact of its author's history (145). Barthes critiques the position of the literary writer, placing it in messianic terms, rejecting the pursuit of the "Author-God," who leaves a trace of "theological" meaning buried like a sacred treasure between the lines of her writing (146). To Barthes, the disembodied authorial entity that literary critics sought to uncover at the root of the text was a phantom. Meaning was not the product of the author and her history—the construction of the meaning of a Text actually occurred in the mind of the reader. A text, he argues, is "a tissue of quotations drawn from the innumerable centres of culture" (146), made up of "multiple writings, drawn from many cultures and entering into mutual relations of dialogue, parody, contestation" (148).

The cohesion of these diverse elements cannot be ascribed to author, whom, Barthes argues, the reader cannot and need not know. The written work is actually granted cohesion and meaning by the mind of the reader, who, for all intents and purposes lacks "history, biography, [or] psychology," and who may be identified as "simply that *someone* who holds together in a single field all the traces by which the written Text is constituted" (148). The "death of the Author," then, becomes the prerequisite for the "birth of the reader" (148). No longer need the reader be

distracted by the person of the author, for the unity of the text is fully formed by the reader herself (Barthes 1977, 148).

I find Barthes' argument compelling, yet I have felt conflicted by it. Barthes contradicts the trope of the composer's centrality to music, an assumption that has followed me throughout my career. I admit the fact of what Barthes says, that a literary work is never transmitted whole into a reader's mind, but rather is a construction of that mind. But I do not believe that the result is an author cast into irrelevance by this cognitive process. When an artwork is seen as transmitting an artist's agency, that force, which guides the action of the artwork, can be attributed to the author, composer, painter, programmer, or writer of that work. The process of the birth of the reader, then, is initiated and shaped by the work, which is itself a process that is set in motion by the application of an artists' agency. If the reader, listener, viewer, or user responds to the work, the maker's agency, I feel, is present, and a source of inertia affecting how the work transforms those whom it encounters. In this sense, the birth of the reader might just be the rebirth of the author, through the invocation of her agency.

The Death of the Composer

Barthes' essay has been located in the stream of poststructuralist thought (Jeongwon and Song 2002, 266). Poststructuralists are not uniform in their beliefs, and there is no central doctrine. There are many characterizations of poststructuralism: that it reveals how humans exhibit a dependency on language where they perceive and encounter objects in their environment (Wenman 2017, 566); that it uncovers "precisely how knowledge becomes possible at any particular time under specific historical conditions" (Harcourt 2007, 18); that it is often conflated with postmodern art movements, and should be regarded as response to structuralism rather than to all aspects of modernism (Ruitenbergh 2018, 693). My interest here is more narrow, simply to explore how poststructuralism comments on music, and what the corresponding implications might be for the integrity of composers' agencies.

Barthes saw a Text and a work as very different, writing in 1977: "The work is a fragment of substance, occupying a part of the space of books (in a library for example), the Text is a methodological field.... the Text is experienced only in an activity of production" (Barthes 2009, 2-3). Barthes' Text may extend beyond a

book—a single text might be the product of a lifetime of writing. The text has no source, in that author’s life is not the origin of its fiction, but rather, the author’s fiction contributes to the work (Barthes 2009, 7). Most important to the current discussion, Barthes likens the text to a musical score, which in his time had increasingly granted the “interpreter” more latitude: “[the interpreter] is called on to be in some sort the co-author of the score, completing it rather than giving it ‘expression.’” (9). In this view, the Text is like a score, wherein the reader collaborates through “playing” the text. To play a text is a pleasurable act, for boredom is the condition of the reader who “cannot produce the text, open it out, set it going” (9).

Must the composer die for the listener to be born? Jeongwon Joe and S. Hoon Song explore this question by examining the parallels between the embrace of chance in music and Barthes’ notion of the Text (Jeongwon and Song 2002, 267). They compare Barthes’ concept of the Text as “a methodological field, experienced only in an activity of production” with Cage’s view that a musical composition is, like life itself, not an object but a process” (267). To Cage, the composition, like the Text, forms in its production, and so compositional control may be minimized, as Cage did by giving up control to chance processes. For example, in his opera series “Europeras,” Cage did not specify notes, instead inviting the singers to choose their favourite arias from other Operas and sing them, a form of quotation that Joe and Song liken to the “tissue of quotations” Barthes declares to be the source of material for the writer’s work (267-268).

While Joe and Song provide an interesting discussion of whether it is the performer or the audience listening, my concern here is more narrowly for the composer. Chance, whether applied through randomization in an algorithmic music context, in the instructions of a textual score like “Europeras,” or even in the unpredictable act where a composer gives a score to a musician whose interpretation adds layers to it, all may be understood as instances where a composer chooses to relinquish unilateral control over musical events. Such a choice is consistent with the collaborative nature of musical activities, where multiple agents (e.g. the composer, musicians, and audience) share responsibility for forming the performance of a composition. To me, the interplay of agencies that emerge in the transmission and performance of a musical work is variable, appearing in different configurations depending on the composer’s aims, the genre, the audience, or the medium of transmission. Cage’s agency

is expressed in the instructions of “Europeras,” just as the composer’s will directs the instrumentalist and the programmer’s agency forms the code. If Barthes’ Author does not make the meaning, I would suggest that the Author facilitates the meaning’s construction. Barthes contends that a Text may be read “without the inscription of the Father” (Barthes 2009, 7), for its meaning need not be attributed to its Author, or indeed anyone else—it is without an origin (7). But that contention has to do with meaning, and my concern is agency. Can the score contain agency? Can the algorithm? If they are processes, and both Barthes and Cage see them thus, there has to be a force that sets them in motion. It is a fundamental property of physics that nothing changes without application of force. Applied to Text and to music (notated or algorithmic), I accept that the meaning may be an internal construction of the Reader or listener’s mind, but I cannot negate the force that set the Text and the Reader into collision, and seeded the sounds that grow from the composers’ garden and smelled so sweet to the listener. That force is agency, and without it, reader and listener might well starve.

2.4 Agency in Music

The existence of interactions of agencies in musical contexts forms a web of social relationships. I will introduce two lenses through which such social relationships are defined. The first, *Werktreue*, privileges the composer, attributing to her agency over all actors in the formation of music through focussing authority within the notated score. This forms a hierarchy, a binary in which the composer is distinguished from other musical agents, through an investiture of authority. The second lens is Christopher Small’s concert of “musicking,” which diverges markedly from *Werktreue*, interpreting music as a shared act that changes relationships among its participants.

2.4.1 *Werktreue* and the Authority of the Score

Origins of *Werktreue* may be traced to the early nineteenth century. The German Romantic author, composer, and critic E.T.A. Hoffmann proposed a new understanding of musical practice: that extra-musical considerations (be they religious, social, or scientific) should no longer guide musical activities; the work itself should guide

all participants in the making of music (Goehr 1989, 55). The term *Werktreue* encapsulates this idea. Translated as “fidelity towards the work” (Stoicescu 2020, 196), *Werktreue* depends upon the separation of composition from performance (Navon 2020, 66). It has been characterized as a “demoting” of performance to a subordinate role in support of the musical work as represented in the form of a score (Hunter 2005, 357). The idea of *WerkTreue* may have seemed alien to the Baroque era’s J.S. Bach, who may have seen his role in humble terms—a devotion to the glorification of God. but certainly was not to Beethoven (Goehr 1989, 56). Indeed, by the 1840s, the centrality of the work to the practice of music was highly advanced. For example, in his famous text on piano performance from 1840, F.-J. Fetis advises students to “render each work according to the thought that created it” and to “contemplate the composer’s work, seize its spirit, and then content himself with rendering it with all the facility of which he is capable, with all the life and sensitivity he can muster, and with as much respect for the productions of others as he would wish for his own” (Hunter 2005, 538). A contemporaneous document by Austrian composer Carl Czerny advises performers of Beethoven’s works that they should “throughout allow no alteration of the composition, no addition and no abbreviation..... For one wants to hear the artwork in its original form, as the Master thought and wrote it” (537-538).

Seen through the lens of *Werktreue*, a musical score becomes a vessel through which a composer may influence an audience, and the act of score-writing reveals her intent, which is the origin of the impact of the work upon the audience experiencing a performance. Of course, as the examples above suggest, this view is rooted in nineteenth-century attitudes, and many composers of the twentieth century reacted against this elevation of the composer, even as many other musical cultures would find the idea unfamiliar. I include this concept here because it is a component that has influenced my own understanding of art music, one that was challenged by my doctoral research. *Werktreue* effectively elevates the agency of the composer above that of other participants in a musical performance, whether that composer is physically present or not. Some Western composers who have influenced me have made claims consistent with *Werktreue*. Igor Stravinsky, for example, asserted his compositional agency in quite forceful terms by stating “I have often said that my music is to be ‘read,’ to be ‘executed,’ but not to be ‘interpreted.’ I will say it still, because I see in it nothing that requires interpretation.” (Stravinsky and Craft 1959, 135; Cook 2003,

204). Schoenberg argued that the performer was unnecessary, except for the fact that concert audiences could not read the printed score itself (Cook 2003, 184). Underlying these perspectives is the granting of a great deal of power to the composer, where they are assumed to exert full agency over the form of the score. The discernment of the composer's intent becomes the object of the music performer and listener, for the score is a document of the composer's intent.

Nicholas Cook suggests that under this perspective, music takes on aspects of a written text, and composition takes on aspects of the act of writing (Cook 2014, 1). To view music so, however, presents an uncomfortable challenge to the performer, since performance is no longer an interpretive act, but one in which the performer may be measured based on how closely he or she expresses the original intent of the composer, found locked within the written score (2). Cook notes that such expectations of performers have prompted the culture surrounding classical music to diverge from many other performance cultures, since the culture of classical music "is regulated in a quite extreme way by the text and those charged with the authority of the text" (1).

While this conception elevating the notated score is associated with the emergence of romanticism in Europe of the late-eighteenth and early-nineteenth centuries (Goehr 1989, 57), and certainly continues to flourish today in some popular and classical music contexts, the centrality of the authoritative, textual score has been challenged many times and by many people. Improvisation, chance music, graphic scores, interactive music, video game music, and indeed musics from around the world that do not share music as notated scores all attribute musical agency in different ways, in some cases even distributing it widely among participants. In algorithmic music, acts of coding may also be seen as an opportunity for the artist to grant agency to others, via sensor inputs and interfaces that admit users input into realtime generative processes. *Werktreue*, then, is a product of a specific period in time and is not universal. Yet, for composers like myself, such centrality may be a compelling notion. This idea is critiqued by Christopher Small in his treatise on "Musicking," a sometimes-scathing exploration of the idea of the classical concert, *Werktreue*, and the latent potential of music to involve people more fully. Since it offers a valuable counterpoint to the narrative of the exalted composer, I now turn to musicking.

2.4.2 Musicking

Christopher Small provides a blistering account of the industry surrounding classical music performance, characterizing the reverence of the composer's score in these critical terms: “[the score] is a sacred object, which is not to be tampered with, whose authority over the actions of all the musicians... is absolute, which commands absolute stillness and silence from those devotees who have assembled to hear it performed” (Small 1998, 122). He characterizes those with authority in the classical performance in religious terms, saying “I intend no sarcasm when I characterize the composer as a kind of prophet, the score as his sacred text, and the conductor as his priest” (87). Under these conditions, the score is the “ultimate centre of power” in the concert hall, and “the symbol of the composer's authority over what is played here and the means by which that authority is exercised” (114). Control of the score, then, becomes control of the people who encounter it in performance, and the composer's agency hangs like a ghost in the air of the concert hall.

Small's stark assessment of the power dynamics in the culture surrounding classical music is predicated on the authority granted to the score by virtue of its status as a kind of sacred text. However, defining music, or its ultimate form, as being equivalent to other forms of written text can be problematic. For example, consider Carolyn Abbate's criticism of the bias of musicologists towards viewing music in textual terms:

Musical sounds are made by labor. And it is in the irreversible experience of playing, singing, or listening that any meanings summoned by music come into being. Retreating to the work displaces that experience, and dissecting the work's technical features or saying what it represents reflects the wish *not* to be transported by the state that the performance has engendered in us. (Carolyn Abbate (Abbate 2004, 505))

Abbate seeks to recognize and elevate the performer, whose action she notes is often reduced to the level of “automatism,” a state where “human bodies [are] wired to notational prescriptions” (508). If we raise the score above other actors in performance, then, Abbate cautions, musicology takes on aspects of hermeneutics: a focus on textual scores which contain a trace of a “secret knowledge” that reveals a modicum of the thought of their composers (510, 533). It is hard to see how such

a music can maintain its audience, where it is entirely built around magnifying the thoughts of others.

Music as a Verb

Small challenges prevailing conceptions of music that treat the score as a sacred text, stating “music is not a thing at all but an activity, something that people do” (Small 1998, 1). The perceived centrality of the textual score to the practice of music, is damaging, Small argues: it implies that musical performance plays no part in the creative process, it renders the experience of music as a one-way form of communication, it constrains the performer’s expressive capacity to the limits inherent in the score, and it separates the musical work from the context in which it is experienced (3-6). All of these effects are corollaries of a widespread assumption Small identifies in which musical meaning is thought to emanate from musical objects rather than from human action (7).

Small introduces the verb “musicking” as a replacement³ for the noun “music,” as he believes that what is most important in musical performances—the essence of their meaning—is how they affect their human participants, and transform their relationships with one another (11). Performance becomes an “encounter between human beings that takes place through the medium of sounds organized in specific ways” (8). We “music” (an action) whenever we take part in a musical performance, in any capacity or role (8).

Barthes’ dismissal of the invisible Author, whose persona is otherwise believed to grant meaning and finality to a text, is echoed in Christopher Small’s characterization of the dead composer. Barthes’ Author haunted the text and appeared to grant it legitimate meaning and an end. Small’s idolized composer presides as a ghost over the classical concert, in an Author-like centrality, the inception point of its emergence. In both critiques, the de facto supremacy of author and composer grant both arguably imaginary entities perceived agency over the substance of their production (book or

3. In using the noun “music” as a verb, Small is echoing a linguistic mechanic used by Nathaniel Mackey and Amiri Baraka (Mackey 1992). Both use the shifting of nouns to verbs and vice versa in their explorations of agency in music and culture, particularly that of African Americans and jazz. Indeed, African-American music and art has long been associated with questions of agency and authorship, a dialogue which seems to have inspired Small in his on critique of the classical concert.

score), and in turn limit its meaning by distracting from the actual meaning, which to Barthes is assembled by each reader and to Small refers to the relationships that form when musickers assemble. Barthes' argument, like Small's, is in part a critique of the tendency of art-related institutions (literary criticism and the classical concert) to enforce a notion that agency over the art object is primarily held by the person whose hand wrote it down, whether that be the author of a text or the composer of a score. The empowerment of the reader, and the encounter with the music listener, requires that they be recognized as active participants in the construction of a work, that their actions actually grant a work its meaning.

2.4.3 Finding Ourselves In Music Technologies

Where is the composer in her music? In the meaning that it evokes? In its distinct form? In the imaginary form of a ghost that haunts the concert hall? In *Werktreue* the composer seemingly owns the score, indeed she might command through it. In musicking, the composer might join the others in a unifying interplay. When digital technologies are used that challenge such roles, how might that change the roles, and change the people who fill them?

Simon Waters observes that, among composers and performers who work within environments that are highly technological, there is a tendency to “celebrate the technological, and to be reductive about (or at least less attentive to) the nature of music as an activity (as practice)” (Waters 2007, 1-2). He associates this tendency with another one, where cultural and historical factors are suppressed in distinguishing the roles of performer, instrument, and environment (2). Waters argues that “musical specialisms”—the distinct musical roles and acts like those of composer, performer, and improviser—have been separated only for a brief time in musical history. Music, he says is a complex, dynamical system, with regards both to its sonic and social organization, and something is lost in these distinctions, that needs to be rediscovered. Waters argues that this process is already underway, as computer technologies are applied to create forms that blend aspects of these roles (as in, for example, laptop performance, DJ culture, and real-time composition and improvisation applications) (2). Arne Eigenfeldt critiques this aspect of Waters' argument (Eigenfeldt 2011, 146), observing that in his experience, not all composers wish to experiment

with fluid improvisational approaches in their practice, although many do.

In my view, there is merit to the reductive argument that Waters critiques. Technology itself plays a role in how we frame our role as musicians, by offering modes of creation that may, or may not, entice the user. The communities of laptop performers require laptops, and DJs require turntables. Yet this is a longstanding issue, and not a phenomenon exclusive to recent history. The composers of the pre-computer era needed pen and score paper as an aid to writing and sharing their scores as well. Their use of these tools, as differentiated from their peers who might have used other tools, creates a separation. The composer is a composer because her self-expressive technology functions well when producing a written score. The pianist who is given the score expresses herself through the piano, a technology that projects sound. I agree with Waters that technologies emerging in recent years blur these borders between roles like composer, performer, and improviser, and that this is an exciting development. I also feel that this process is part of a larger arc of evolving human relationships with technology that invites us to explore new roles, whether blending previous roles or innovating new ones. Music-related technologies have always extended human roles and sometimes introduced new ones (e.g. those of the piano virtuoso, music theoretician, VST plugin designer, or algorithmic composer).

Waters provides an wonderful image of the connection between human, instrument, and environment—how they contact and merge with one another, something that might be overlooked when performance is conceptualized as a heterogeneous interaction between performer, instrument, and environment. He first asks what is lost in accepting these separations, then answers:

One victim is an important ambiguity, the fragility of the performer-instrument articulation—the specificity of an individual’s ‘touch’—which results not only from the physiology of the player, but the complex feedback into that player’s body of vibrating materials, air, room, and the physiological adaptations and adjustments in that body and its ‘software’ which themselves feed back into the vibrating complex of instrument and room. Think of the extension of an instrument’s capacity back into the body and physicality of the performer (the tube of the flute, which I was taught at age eight to think of as starting at the diaphragm and extending

into the room). Think of the multiple affordances of the instrument: Although the modern Boehm clarinet is ‘designed’ as an equal temperament device, this cultural expectation is carried as much in the performer’s body as in the acoustic system of the physical object, and in the hands of a South Indian musician the same device affords entirely idiomatic delivery of a music which is subject to entirely different principles (of pitch subdivision and much else). (Waters 2007, 2-3)

Waters observes that a performer and instrument are connected by “touch,” a wedding of the player’s physiological interactions with the instrument *and* the environment around them. When a performer plays through a clarinet, he says, the limitations and creations that ensue are not mostly concentrated in the instrument, but depend on the player’s body, and the mental “algorithms” that animate it, and in the relationships between these and the environment (with its social and acoustic dimensions) (3). When we engage with computers, we may attempt to enforce a similar distinction of roles—it is an instrument when we perform on it, an environment when we build something in it, and a performer when we code it to create something automatically (4). In reality, the three are not so easily distinguished in the computer, Waters says. He introduces the term “Performance Ecosystem” as an enfolding of these three concepts, and a mitigation of their distinction (4-5).

Waters’ metaphor is apt. A biological ecosystem encompasses a diversity of elements, species of plant, animal, fungus, and microbe, predators and prey, all acting within an environment that contains them even as it is them⁴. To treat music reductively as a multiplicity of distinct roles just might obscure the whole, the aspects of music exterior to the roles that are found in interaction. What is a bee without the flower, or a bird without sky? They exist, and can be described. But the networks that connect them to the other organisms and things are also meaningful. The ecosystem behaves, shifts, grows, and contracts, its shape is formed by a mass

4. Waters’ metaphor of the ecosystem also evokes Deleuze and Guattari’s idea of “assemblages,” where a multiplicity of entities, physical, corporeal, or conceptual, come together for a time in the realization of a function (Parr 2010). It also relates to the idea of networked music, where computer systems sustain collaborative music-making across performers who are physically-remote from one another, often using low-latency internet connections (Rottondi et al. 2016). While both of these concepts are beyond the scope of this thesis, they both imply a “coming together” of disparate forces, a function which, in empathetic terms, I find inspiring.

of interactions that emerge from individual species. So, too, is the case of musicking, where relationships are the product of interaction that is facilitated by shared engagement in musicking.

2.5 Relevant Artists

Many composers have a conception of agency that is more nuanced than implied by the dictates of *Werktreue*, or by the equalitarian critique of the classical concert music institution by Small. I find that these two viewpoints leave a fairly wide territory in-between them where composers can exert agency over their scores, but still admit the agency of others, whether it originates from listener, musician, or some other entity that is welcomed into the process of composition. As such, I will next introduce some significant composers who claim some of this territory.

2.5.1 Lejaren Hiller and Leonard Isaacson

Some of the earliest work in computer-based algorithmic music included Lejaren Hiller and Leonard Isaacson. They are credited with designing the first automated application for composing music, a program which was used to compose the “Illiac Suite,” a notated score for string quartet (Liu and Ting 2016, 1). While the programmers’ stated aims were of a technical nature—different movements demonstrated that computers could compose autonomously, help composers, or innovate in musical form (Pearce, Meredith, and Wiggins 2002, 120)—it does represent the first recorded instance where individuals chose to deploy a digital computer to compose a theme, rather than compose it themselves. As such, the question as to the consequences of relocating composer’s agency from the language of notation to that of code begins here, and this composition serves as the inception point of the central problem that this thesis addresses. I will discuss this composition in more detail in Section 2.7.2.

2.5.2 John Cage

The influence of aleatoric composers like John Cage on conversations of agency is significant, especially his well known composition “4.33.” Its score reads thus: “Tacet. For any instrument or instruments. The piece is in three movements: .30, 2.23, and

1.40” (Davies 1997, 448). Does the score contain music? Not in the form of notes, to be sure. But Cage did not intend the score to be empty, saying that the piece “becomes in performance the sounds of the environment.” He elaborates: “the essential meaning of silence is the giving up of intention” (Kostelanetz 2003, 198), which in 4’33” allows the ambient sounds from outside and within the performance space to appear. Yet the work itself, like other musical scores, does direct the performer to carry out a task, specifying—to the second, no less—the durations during which the instrument(s) are silent, and the audience listens to what was always present, but never heard. This situation set up in the score is an intent. Cage’s nonintention is not the absence of intention.

He describes two ways to shift intention to nonintention: “musicircus,” in which many intentions collide to form an absence of intention (158) and “music of contingency,” improvisation with instruments that introduce a discontinuity between cause and effect” (64). Musicircus is an apt term that describes a fear motivating my early algorithmic music projects described in this dissertation: that without a composer’s authoritative agency guiding the form and effect of unfolding music, its meaning would dissolve in a cacophony of competing agencies. Cage challenges this notion, that a singular vision should be projected through music, saying of audiences:

“They should listen.... They’re convinced that it’s a vehicle for pushing the ideas of one person out of his head into somebody else’s head, along with—in a good German situation—his feelings, in a marriage that’s called the marriage of Form and Content. That situation is, from my point of view, absolutely alarming” (249-250).

For Cage, focusing the meaning of a work introduces the ego, carrying with it likes and dislikes that obscure the qualities of its sound (Kostelanetz 2003). Though he had been taught, as many other young composers had, that music is effectively an expression of the self (i.e. the ego), he observed that music did not carry meaning intact—sad music was interpreted as happy, no-one could truly understand the other through hearing the work, forming a “Tower of Babel”-like situation (230). Yet Cage still considered himself a composer, only he has changed the location of his responsibility from choosing events to asking questions (228). His intent is not to express, but to “make something that can be used by the person who finds it expressive” (230).

Here, Cage, like Barthes in “The Death of the Author,” places the formation of the meaning of a work to the receiver. This is a comforting thought to me. I have feared that the product of works where the composer or programmer relinquishes too much control to other agents could lead to a cacophony—many voices speaking but few listening to one another, and the whole consequently lacking a focus. When applied to algorithmic composition governed by multiple agencies, the idea that meaning is a construction of each agent, rather than an ideal imposed upon the work by a authoritarian composer, suggests that the cacophony I feared may be avoided, replaced with a chorus of subjects constructing their own meanings in a shared space. This idea is compelling to me, dissolving the distance between composition of music in traditional notation (which in Western tonal music may be characterized as linear, and under the control of the composer) versus that occurring in code, with entry points for variation that might be governed by composer, programmer, performer, user, chance, or other environmental factors which the system can sense. If the listener in each case is in their own mental domain, constructing their own meanings, why does interactivity lead to a cacophony? After all, in both linear and dynamic music, the music is experienced as a temporally-progressing sequence of events. I am not ready to dismiss the ego, as I believe that music can form connections between people, and that if likes and dislikes obscure the qualities of the sounds, that is a form of meaning. Nevertheless, Cage’s conceptualization of intent, where choice gives way to questioning, provides a starting point for describing the relinquishing of agency that happens when algorithms automate composition.

I will return to Cage’s work in Section 2.7.3, where I introduce his relationship with algorithmic music.

2.5.3 Iannis Xenakis

Among the earliest composers to embrace the use of programmable computers as a source of compositional material and variation is Iannis Xenakis. I introduce him here for two reasons. His conceptualization of “stochastic music,” which I describe in more detail in Section 2.7.4, shifts control of some aspects of musical event structure from the medium of notation to that of code, which I will frame as a transformation of composerly agency. What I want to briefly discuss here is Xenakis’ choice to embrace

acts of programming in the definition of musical elements.

Xenakis had already developed probabilistic systems to determine values of musical variables in the works “Metastasis” and “Pithoprakta” (Serra 1993, 237). The latter work is particularly interesting, and while a full accounting of the work is beyond the scope of this dissertation, some of its essential features are relevant. First, the score is a representation of only one possible form of the piece—variation in musical events derives from Xenakis’ implementation of a probabilistic model derived from Maxwell and Boltzmann’s kinetic theory of gases (Arsenault 2000, 14). Each of the 48 string players performs an independent line governed by “speeds,” measures of pitch changes that are determined by a function forming a Gaussian distribution⁵ (15) rather than by his own choice of pitch.

Why would he choose to model music structure on the behaviour of air molecules, rather than his own intent expressed via conventional notational methods? Perhaps his background played a role. An autodidact composer, he was also trained in Greek Philosophy and engineering, and studied architecture (Reynolds 2003, 4-5). In his early career, he acted as architect and composer, a parallel evolution which his associate Roger Reynolds describes as “a conjunction [which] was a signal illustration of the rare composite of forces which intersected in him: objective and mathematically principled knowledge conjoined with a raw emotional directness that could be at times alarming” (5). He wrote articles revealing mathematical principles as opportunities for musical invention, but neither his articles nor his scores found large appreciative audiences, but rather were more commonly greeted with incomprehension.

This brief and quite incomplete characterization of Xenakis’ life and work nevertheless illustrates two themes occurring in Xenakis’ work that are fundamental to this dissertation. First, Xenakis’ education is rooted in interdisciplinarity, as his education in engineering cannot easily be separated from his compositional methods. Second, his use of the computer to facilitate the construction of his works can be understood as an early instance of a form of computer music where traditional facets of composition

5. When a Gaussian distribution is found in nature, it exhibits a particular distribution of values. These values are anchored by two variables that may be assigned arbitrarily: the mean, or central value of the variables, and the standard deviation, a value of some specific distance away from the mean. In a normal distribution, about 68% of values would fall within one standard deviation, 95% fall within two, and 99.7% within three. Values outside of this range are rare, and may be considered negligible (Arsenault 2000, 26-27).

are not constraining the form of a piece. Hiller and Isaacson programmed their mainframe to write a string quartet. Whatever technical innovations they were exploring, they had not transcended classical music itself, instead producing through code that Mozart and Haydn themselves might have identified as a quartet, by arrangement if not language. Xenakis' music did not come from the eighteenth, nineteenth or twentieth centuries. It came from him, and the rules and visions informing the design of his algorithms.

I will return to Xenakis in Section 2.7.4, where I discuss stochastic processes.

2.5.4 Roy Ascott

New media artist Roy Ascott has been central to discussions of cybernetics which directly inform my ideas of agency. “Cybernetics”, according to Ascott, has been described as “the study of control and communication in animal and machine” (Ascott 1968, 106). The definition is consistent with the earliest conceptions of the cyborg by Manfred Clynes and Nathan Kline, who in 1960 introduced the cyborg as a self-regulating feedback loop between human and machine (Clynes and Kline 1995, 27). Ascott's application of the cyborg is informed by a recognition that the art of his time was changing. He said of older forms of art that they were “deterministic, concerned with discrete objects, with things and fixed relationships.” On the other hand, he identifies in the art of his time the qualities of being a “process” and “open-ended” (Ascott 1968, 106). He observes a “shift of human interest” from objects and products to processes, systems, and the events through which they create a product (106). The theories of cybernetics had a long history in his work. In 1961, he was asked to form the “Ground Course,” a cybernetics-themed course for artists and designers at the Ealing Art School (Quaranta 2013, 52).

Ascott uses the term “telematic art” to explore the nature of meaning in a creative environment that, in its use of computer technology and the interaction models that ensue, is quite unlike the world of the past (Ascott 1990, 241). “Telematics” refers broadly to computer-mediated communications networks, and new models of interactivity among and between human beings and artificial systems. Telematic art, made using such interactive, networked media, involves its audience not as mere observers but as participants, and, in an echo of Barthes, meaning is not received by

its audience but emerges from the interactions between observer and system (Ascott 1990, 241). The culture sprouting around such innovations is, broadly, one sustaining contact between participants. Ascott describes it thus: “Telematic culture means, in short, that we do not think, see, or feel in isolation. Creativity is shared, authorship is distributed, but not in a way that denies the individual her authenticity or power of self-creation...” (243).

In Section 2.8.6, I will introduce Ascott’s work “Aspects of Gaia,” which illustrates Ascott’s application of telematics and cybernetics in art.

2.5.5 Hans Zimmer

Film composer Hans Zimmer has composed or collaborated on the composition of soundtracks for some notable films, including those for “Inception” (2010), “Interstellar” (2014), and “Blade Runner 2049” (2017). I believe mention of Zimmer belongs here for two reasons: his embrace of collaboration, and of digital composing technologies.

Zimmer’s collaborative approach to scoring a film speaks to a model where multiple agencies, the composer’s and otherwise, govern the form of a score. Zimmer has developed a corporate organization around himself called “Remote Control.” Recent advertising for the company boasts that “clients have access to over a dozen composers and music editors,” positioning the collaborative nature of the process as a main selling point (Kmet 2018, 1). An early work of Zimmer’s produced under the banner of Remote Control was the film “Pirates of the Caribbean,” which I understand as an illustration of the complexities of composers sharing agency over a single musical work. Zimmer was asked by producer Jerry Bruckheimer to produce a score on a tight timeline after composer Alan Silvestri departed the project (8). While Zimmer takes credit for writing much of the thematic material for the film, the composer’s credit is assigned to Klaus Badelt, whom Zimmer initially assigned the task of completing the score when it was first offered due to his time being committed to another project (7). In fact, although initial themes were devised primarily by Zimmer and Badelt, the tight deadline for delivery required contributions from

many at Remote Control—the iMDB page lists eight other composers as having written “addition music”⁶, as well as numerous orchestrators, music editors, and music programmers. The score has been criticized, despite the film’s popularity. Writing on the FilmTracks.com website, Christian Clemmensen has called the soundtrack a “monumentally disappointing mess of a score.... consider the days when a single man would write, orchestrate, conduct, and produce a score. Now imagine two-dozen people trying to do the same thing all at once on computers, and the product is a useless, meandering collection of stock action cues with few cohesive elements of any significance” (Clemmensen 2007).

On the surface, positioning the score as a failure due to the lack of a single centralized voice behind it might seem quite appropriate, and I feel it reflects the attitude I outlined earlier where the composer is seen as privileged, in total control of the score, and as the unitary source of its expression. However, in more recent projects, Zimmer has maintained this approach with more critical acclaim. The score for the 2008 Christopher Nolan film “The Dark Knight” was a joint effort by Zimmer and James Newton Howard, a rare instance of two highly-regarded film composers working together on a single score (Hexel 2016). In an interview about the score, Zimmer states:

“Once you get the job of being the film composer the moment you say yes to the movie and you sit down to write it becomes the antithesis of creativity in a funny way because of the pressure of having the burden of being the last person in line that could make or break this movie. Just the idea of very often having to fit within the framework of something else is not what we imagine to be the best creative playing field, which is freedom and imagination and chaos to a certain degree. But that’s where the relationships come in. So if you make it about the human thing, if you make it about ‘hey I actually think this director has an interesting life’ or his soul adds something to the way I get to spend my days.... At the end of the day that is what it comes down to, these are our lives and the seconds of our lives are ticking away and with out getting all morbid

6. Additional composers who worked on “Pirates of the Carribbean” included Ramin Djawadi, Jim Dooley, Craig Eastman, Nick Glennie-Smith, Steve Jablonsky, James McKee Smith, Blake Neely, and Geoff Zanelli.

here, the best of quality of life you can have, for me anyways, is to have an interesting conversation, to have a connection with somebody.”—Hans Zimmer⁷

Zimmer here contradicts the view of the composer as a solitary creative with a singular vision, emphasizing instead the value of collaboration and conversation in realizing music. That his and Howards score to “The Dark Knight,” and indeed many of Zimmer’s other recent scores, received favourable reception from critics, only supports the contention that visionary art need not emerge from a singular vision. Rather, it suggests that in an interplay of agencies, meaningful art may emerge, that a blending of perspectives from a dialogue of multiple minds can succeed. The above review of the “Pirates of the Carribean” score laments the days when a single individual served as the unitary lens through which film music was transmitted, a view that I have shared at times. Yet Zimmer’s collaborative impulses do not necessarily cheapen his work, dilute it, or cause it to fall out of focus. It is that spirit, I feel, that needs to inform the creator of algorithmic music, where they facilitate and join a conversation. I believe that is a direction that music is destined to approach as the constriction of music is relocated from score to algorithm, and it is a conversation worth having.

Zimmer’s collaborative approach to film scoring may be linked to his embrace of technology as a facilitator of his and others’ self-expression, such as his use of the digital audio workstation Cubase (which I myself used when composing “We, The Artificial”), creation of custom plugins and sample libraries for his studio work, and interest in creating music-related apps. He has crafted music for nonlinear, software-bound musical forms, including video games like *Modern Warfare 2* (Fritsch 2013, 11), and the an interactive realization of his score for the film “Inception” as mobile phone app (Plans 2017, 57).

Both the collaborative methods and technological environment by which Zimmer distinguishes himself as a composer have been criticized. His practice challenges the idea of the composer as a sole author, and raises issues of attribution which are sometimes framed as a straw man where Zimmer is cast as a mere delegator (Kmet

7. Taken from a transcript of an interview with Zimmer and Howard by Jason Bentley. The full text can be found here: <https://www.kcrw.com/music/shows/morning-becomes-eclectic/hans-zimmer-and-james-newton-howard>

2018, 2). Other critics react to his porous concept of boundaries of style. His use of synthesizers, ethnic and pop music instrumentation, and deployment of a harmonic language that largely ignores avant-garde developments outside of the tonal Common Practice have led to accusations that his work is derivative and simplistic (Lehman 2016, 32).

In fact, Zimmer's process is more complex, both in its collaborative methods and in his use of musical and technological tools. He typically begins a scoring project by composing a twenty-minute suite, comprising all of his intended themes for the film. Once the suite is approved by the director, it is divided and applied to the film as a temporary score. At this point the other composers enter the project and flesh out the themes, arranging them under Zimmer's supervision (Kmet 2018, 3). Whether this is done for speed or for artistic invigoration, this collaborative framework does challenge established notions of authorship in film music, blurring the line between composer and producer, and granting greater agency over the form of a score to film editors (12).

Despite these and other criticisms of Zimmer, I think Zimmer's collaboration-based model of composing is the challenge to authorship that music needs today. Arguably, Zimmer has reached a high level of critical and popular success because his prodigious capacities as a composer of accessible themes have been extended by technologies with which he is able to exploit effectively. Zimmer's collaborative model of shared compositional agency is sustained by his adept deployment of technologies which enable new methods of solo and shared composition ⁸ . A result is that his most collaborative film scores are multi-agentic creations, overturning the common notion of the solitary film composer. Thus, I find that Zimmer's approach affirms the potential of interactive algorithmic music to reach the heights of significant film composers, and this is a gift to the algorithmic music composer who, like myself, looks towards a near horizon where the interactive score can express a subtlety, depth, and complexity of the finest film scores, even as it sustains a conversation between the agents around whom it forms.

8. Zimmer talks about his studio and use of Cubase in a video here: <https://youtu.be/141F6hfbyhc>

2.6 Agency Through and Of Machines

I will introduce three topics here that relate to how agency can be conceptualized as a attribute of a machine in the discipline of computer science: software agents and frameworks, human/computer interactions, and autonomous agency.

2.6.1 Software Agents and Frameworks

There are many conceptions of agency in computer science, and the word “agent” has been described as an “umbrella term” (Nwana 1996, 6). Software agents emerged from the field of Distributed Artificial Intelligence (DAI), (2), which innovated through designing systems of entities that collaborate. Such systems have two defining characteristics: They are designed to solve a problem, and solve it by deploying subsystems (the “agents”) that work autonomously but collaboratively to that end (Ponomarev and Voronkov 2017, 5). This innovative, multi-agent approach to computational problem-solving has been attributed to Carl Hewitt (Nwana 1996, 2), who, in a 1977 article, (Hewitt 1977) proposed the multi-agent system as an alternative to systems modelled upon a single human being thinking and acting to resolve all aspects of a given problem (323-324). Hewitt explains that in the prevailing model, the system is conceived as an expert, containing all capacities necessary to solve the problem for which it was produced. His alternative system forms a “society of experts,” in which objects he called “actors” (a term now synonymous with “agents”) assume different roles based on their varied intrinsic capabilities (325). The solutions to problems are beyond the capacity of the component agents, but emerge from cooperation between them. Agents form a quasi-social environment, which grants the group greater power to solve problems than any individual agent might possess.

Properties of a Computational Agent

Drawing upon Hewitt’s ideas, Hyacinth S. Nwana lists several capabilities found in an agent, including abilities to cooperate, learn, and act autonomously (Nwana 1996, 7). If an application does not exhibit at least two of these characteristics, Nwana declares that the application is not truly an agent (7). A contemporaneous writer, James Ingham, offers a “working definition” of an agent, describing it as:

An entity having the property of controlling its own action independent of other entities unless it desires to communicate with other entities. It typically does not refer to a second party unless it possesses insufficient knowledge to perform that particular action itself (Ingham 1997, 11).

Another definition is provided by Franklin and Graesser 1996, who emphasize the agent’s relationship with its environment:

An autonomous agent is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future (Franklin and Graesser 1996, 25).

Franklin and Graesser situate agents along a continuum. On one extreme, there are humans (and some animals) possessing “multiple, conflicting drives, multiple senses, multiple possible actions, and complex sophisticated control structures” (25). At the other extreme objects like the thermostat, which Franklin and Graesser describe as having “one or two senses, a single action, and an absurdly simple control structure” (25). They argue that their definition of an agent encompasses both of these extreme cases—the divergence between human and thermostat is one of degree, not substance. I find Franklin and Graesser’s to be the most compelling technical definition of an agent, in that it situates humans on a continuum along with simpler systems, including some of my own coded projects. I feel that it can be helpful to assume an equivalency between human and computer as regards their identities as agents, especially where those agencies are set to interact.

The three preceding definitions were advanced in an early stage in the development of software agents, but are largely consistent with current definitions (see, for example, Mostafa et al.2017 and Ponomarey and Voronkov 2017). While the idea of the agent continues to be used, a more recent innovation, the software framework, provides an alternate idea of agency in code.

Software Frameworks

A software framework is a tool which supports the the development of custom software applications (Fayad and Schmidt 1997, 32). They can contain predefined,

generic components which the user of the framework can deploy to achieve their goals (Fayad and Schmidt 1997, 33). Frameworks govern communications between the user-coded objects in response to events, a programming principle called “inversion of control” (34) that simplifies the users’ design process. Frameworks provide an alternate model of agency to bots and software agents. The designer of an application codes extensions of the skeleton that is the framework, but when the execution of the application, the flow of control is governed by the framework.

Software frameworks blend the agencies of programmers. The frameworks user, who is designing an application, gains the benefit of the function of the framework, but in turn has to conform to its structure. An example is Unity, which is a framework for creating video games and interactive media. Unity comes with a large number of capabilities that can be used to create an application, such as methods for manipulating graphics and audio, receiving input from game controllers and other sensors, and communicating over networks. This extends the capability of a user, who does not need to code such capabilities from the ground up. In my own use of Unity to build the “Shards of Memory” iPhone app (see Chapter 7), the framework gave me access to sensors, audio outputs, and camera streams, which I extended through my own scripts that formed the app’s generative engine, and governed its behaviour. While such an environment might seem liberating by in granting access to a range of capabilities, frameworks necessarily impose limits. Through inversion of control, the framework enforces conventions that may, consciously or unconsciously, prompt the user to develop in certain directions, while closing off others.

As debates within the field of User Experience (UX) and Human-Computer Interface (HCI) design continually remind us, interface design carries what new media artist Rokeby calls “a social responsibility” (Rokeby 1998, 1). I agree, and see this responsibility in terms of a interplay, or even struggle, between agencies. Framework designers and framework users may have divergent aims, and it is perhaps a implicit risk that one may suppress the other. I identify this dynamic in interactive works such as Brian Eno’s “Bloom” (Section 2.8.7) and my own “Moonloops” app (Chapter 6). Where such agentic interplay carries the potential to suppress or deceive, even to a small degree, design of a system can and should be considered in such terms⁹ .

9. As an interdisciplinary artist I consider the experience of users of my apps, but in a different way than I consider the listeners of my scores, since where an app is interactive I believe it sets up,

Agents Versus Non-agents

Not all algorithms are agents. Franklin and Graesser give the examples of a simple spell checker or a payroll program (Franklin and Graesser 1996, 26). These programs are not agents because their behaviour is utterly constrained—they do not learn, and are unaware of their environment beyond inputs to which they respond with absolute consistency. True agents are aware, reactive to the environment, and unpredictable.

Here, I can imagine Sartre raising an eyebrow. This definition of an agent does not only distance it from the tool-like calculator or spell checker that have no capacity to surprise us with its mathematically-guided determinism. Such devices are like rocks, they cannot act except according to the dictates of their code. This inertia is at the core of Sartre’s being-in-itself. Franklin and Graesser distinguish a software agent from other programmatic types through capacities that might cast doubt on its inherent facticity, and permit it to act with a modicum of choice. I have wondered if this quality of a software agent might be the beginning of an algorithm becoming a being-for-itself. After all, if a software agent can assume humanlike qualities—those that Sartre ascribes to a being-for-itself—is it valid to deny its agency, or its capacity to hold an autonomous will? Of course, this question is not answered in this dissertation, but it is raised because I believe it points to a distinction between composers and coders: the notation medium of the former cannot resist, but that of the latter might.

2.6.2 Computer-borne Agency in Human Social Contexts

To me, one of the most interesting aspects of software agents is their integration into human societies. Software agents can collaborate, form an interface or mediate, move in physical or data space, and act across networks (Lebeuf et al., n.d., 2), all of which can be seen as the acts of social beings. Can we grant a software-borne agency such a status?

to a greater or lesser degree, a kind of social relationship. The fields of UX and HCI explore such issues, but my intent in this thesis is not focussed on fields, and they are outside of the scope of this research.

Video Games and The Implicit Human

A rich and significant field in which interactions between humans and machines are a central concern is that of video games. A video game can represent not just fictitious objects and characters; it can also learn about its users and, through acting upon their in-game avatars, become a site where human social relationships are carried out. In my future research, I aim to explore the social applications of adaptive music more thoroughly, but, even in the work presented in this thesis, I have been mindful of the tendency of users to project themselves into a game world, and my work “Shards of Memory” includes an avatar construct to mediate between the user and installation’s content. For this reason, I will introduce some ideas from researchers who study video games, beginning with the blurring of real and virtual that may occur when one is immersed in such as experience.

Tilo Hartmann explores the behaviour of video game players when they engage in simulations of violent acts (Hartmann 2017). His “Moral Disengagement in Violent Videogames” model can be summarized through the following four propositions (2):

- Players intuitively perceive video game characters as social beings;
- Processes of socialization deem social beings to be worthy of proper moral treatment;
- Mistreating a video game character triggers “empathetic distress,” with the associated guilt diminishing enjoyment; and
- Video games are forms of entertainment designed to be enjoyed, and so the designers embed elements that provoke moral disengagement, justifying the player’s actions through absolving them of guilt, so permitting players to enjoy the violence.

I appreciate Hartmann’s analysis, there is an argument to be made that it can be extended from the virtual environment of a video game to situations in the physical world, where we encounter machines that exhibit human characteristics and behaviours resembling social interactions. Psychologists have studied this human tendency since the 1940s, when Heider and Simmel 1944 examined how humans interpret nonhuman behaviour. They presented participants in their study an animation of two

triangles and a circle moving in a 2D-space, and asked them to describe the objects' actions. Almost all participants described the objects' actions as if they were humans in a social interaction (Heider and Simmel 1944, 246). One participant sees in the shapes' relative movements an argument between two men, another interprets the circle and one of the triangles as lovers (246-247). This research, conducted long before video games were invented, suggests that Hartmann's analysis relates to an essential facet of humans' perception of the world.

This misattribution is called "anthropomorphizing," and, according to Heberlein and Adolphs, it occurs where humans "attribute social meanings to stimuli that are not social, such as computers or clouds, presumably based on cues that signal the presence of agency or emotion" (Heberlein and Adolphs 2004, 7490). The phenomenon is a profound influence on human relationships with machines, appearing when we name ships as people or cry when a beloved video game character's virtual life ends.

Controlling The Machine

The blurring of social boundaries between human and machine has been exploited by storytellers for at least a century. Examples include HAL, the infamous artificial intelligence in Stanley Kubrick's "2001: A Space Odyssey," whose calm, helpful personality masked malevolent intent (Kubrick 1968). There is the trope of the emotionless and spiritless artificial human, seen in the persona of the Tin Man in "The Wizard of Oz," who desires a heart to complete his humanity. One of the earliest depictions of an artificial being in film is in the mechanical dystopia of Fritz Lang's "Metropolis," where a hero of the film, Maria, is impersonated by a robotic facsimile crafted by an evil scientist, deployed with the intent of destroying Maria's reputation, and undermining her advocacy on behalf of the workers toiling the bowels of the city. The robot ultimately riles the proletarian workers into a destructive mob (Allehaibi 2019, 160). The robot displays what contemporary doctors would have termed "hysteria," where it dances seductively, incites violence, and laughs during its own execution (Kang 2018, 39). Hysteria can be understood as the pathologization of the womb, seen in writings ranging from Hippocrates of Ancient Greece to Sigmund Freud (Krasny 2020, 128). A woman with hysteria was viewed as out-of-control of her body and its reproductive capacity, which necessitated the "shoring up" of patriarchal power over generation (128-129). The hysteria of the robotic doppelganger of

Maria in “Metropolis” manifests just such a unruly nature—where the real Maria is kind and helps her fellow dwellers of Metropolis, the robot is disordered, destructive, and dangerous, and will not submit to the patriarchal governor of the elite, upper domain of the city (Kang 2018, 39). It is an instance of the feminization of the machine, an undesirable framing of femininity in Victorian England where women were characterized as mechanistic, apt only to engage in repetitive and automatic tasks that required no substantial will or autonomy (35). Men, in contrast, were framed as engines, creative and progressive, and, lacking wombs, they were less prone to the disorders of irrationality. The hysterical “woman-machine” reflects a fear of the uncontrollable, a threat to the sanctity of human invention. It also reflects a gross misconception of women, one averse to acknowledging female agency or thought.

The machine itself is subject to just such an anxiety-fuelled diminution today—we may fear its becoming detached from our direct control, for we know that the range of its autonomous action may have no appreciable limits. The Victorian woman who was constrained by a medical diagnosis of hysteria was not being aided by the doctor. Rather, the doctor was aiding the patriarchy that sought to suppress her agency. When a machine gains autonomy, and makes something which its engineers did not anticipate, might we not similarly be frightened? And might that fear not be as unfounded as that of the supposedly-hysterical Victorian woman? In “Metropolis,” the fear is justified, but that, too, is a fiction. The fear of autonomous computer programs, including the variety used to produce algorithmic music, might be tempered were we simply to ask where the fear came from. The machines themselves, or our own fear of losing control? Though the answer remains in the future, we must be careful in our present uncertainty. Though autonomous agents may spark unease, to reject them outright might yet be foolish, as we extinguish a spark that might have opened paths to new knowledge or human expression.

The presentation of these machines in film and literature speaks to an unease in humans’ encounters with these machines—they are like humans, but incomplete, and often this artificiality accompanies tragedy. Isaac Asimov’s “I, Robot” series of novels famously introduces his “three laws of robotics,” a foresightful codification of principles to guide artificial beings’ interactions with humanity and one another. In the books they are said to be designed to protect humans from the assumed threat posed by autonomous machines, adding depth to literary treatments of social machines that

generally saw them as an unconstrained threat (Balkin 2017, 1218). These three laws, forbidding machines to harm humans, other robots, and themselves (in that order of priority), were Asimov’s response to rising concern that our autonomous creations could eventually turn on us, a belief that Asimov termed the “Frankenstein Complex” (1218). While robotics researchers have argued that implementation of the laws are not practically feasible (Kaminka et al. 2017, 344), they have sparked a debate over the use of machines, one that becomes more important as machines are designed to slip ever more effortlessly into human social roles.

Cautions Against Social Equivalence

There are numerous cases where the design of agents to act as social beings has been motivated by deceit. An early example is Joseph Weizenbaum’s “ELIZA,” which could emulate a human conversational partner (Candello and Pinhanez 2016, 2) to deceive its users into believing that they were talking with a human, Rogerian psychotherapist (Candello and Pinhanez 2016, 3; Shum, He, and Li 2018, 11). Possessing extremely limited knowledge of psychotherapy, ELIZA used pattern-matching algorithms to analyze human conversational partners’ typed input and constructed responses based upon scripts that were triggered by keywords in the input (Shah et al. 2016, 280-281). If ELIZA’s capabilities improved the wellbeing of any of its conversational partners, however, the effect would seem to be incidental. Yet some of its users took its “advice” seriously, even after learning that it was actually a machine (Candello and Pinhanez 2016, 2), an interpretation that, while it might seem laughable today, only provides evidence of the capacity of software desires to play on others’ tendency to see their environment through a social lens.

ELIZA has been described as the first program to pass the “Turing test” (3) Alan Turing proposed the test he called the “imitation game” in 1950, as a method of evaluating the intelligence of a artificial intelligence system (Turing 1950). He argued that the answer to the question “can a machine think?” may be found in determining whether a machine’s behaviour appears indistinguishable from a human’s (433), While his Turing test (as it became popularly known) may be too simplistic to thoroughly address the question of what constitutes artificial intelligence, it does frame it as a problem of distinguishing human and machine. Recent technological advances suggest the value of such a view. For example, Ferrara et al. observe that social bots acting in

online social media forums can influence human opinions and beliefs by “infiltrating a population of unaware humans and manipulating them to affect their perception of reality, with unpredictable results” (Ferrara et al. 2016, 99). They discuss Twitter bots, algorithms which create accounts on the Twitter social network platform and emulate human posting patterns. Of course, such bots are not human, they are tools deployed by hackers for variety of purposes, including the spread of misinformation for political purposes or to tamper with real people’s reputations (99-100).

These potentially-destructive phenomena are successful partly because such agents have effectively assumed human identities on social media platforms. The successful construction of illusory human identities *en masse* may grant a hacker an immense population of socially-influential entry points into political and social discourse, much more than the hacker would be able to muster were he or she limited to a single, manually-curated social media account. The hacker who wields social bots is essentially magnifying his or her singular perspective by illicitly multiplying his or her identity in the form of bots can only agree. The bots do this by assuming human status, and their unwitting co-habitants in social media may then tacitly accept their action as human, granting them the illusion of possessing intrinsic agency that may actually emanate from their concealed human designer.

2.6.3 Relating To And Through Machines

Jack M. Balkin has coined the term “homunculus fallacy” to refer to humans’ manner of relating to algorithms, robots, and artificial agents in general. (Balkin 2017, 1223) The homunculus fallacy reflects an underlying assumption Balkin believes we may form (consciously or unconsciously) whereby we interact with an algorithmic agent as if it possesses a little human within itself (i.e. the “homunculus”) that controls its behaviour. This homunculus, we might think, has good or bad intent, and may choose to do good or bad things. The flare of anger I sometimes direct towards Siri when it misinterprets a perfectly-enunciated instruction is evidence of the homunculus fallacy at work. I forget for a moment that Siri essentially cares nothing for such anger because Siri is a software construct and does not possess an innate humanlike capacity to feel corrective emotions of shame or fear that my angry words might raise when directed towards a person.

Balkin cautions that displacing responsibility for an algorithm’s actions into the application itself and not towards its designer can obscure significant power dynamics that exist between all people who encounter one another through a machine. He writes, “the effects of robotics are always about the relationships of power between human beings or groups of human beings” (Balkin 2017, 1225). He gives the example of an algorithm that was designed to pick beauty pageant contestants. An account of the event is provided by the “Guardian” newspaper (Levin 2016). The 2016 beauty contest¹⁰ was intended to be the first such competition to be judged by algorithms using objective criteria such as degree of facial symmetry and appearance of wrinkles. It drew submissions of photos from approximately six thousand contestants representing over one hundred countries. Once the forty-four winners selected by the algorithm were revealed, the winners list was observed to be overwhelmingly caucasian, despite the diversity of body forms and skin tones among the pool of contestants. The algorithm was actually accused of racism—including in the article in the “Guardian” newspaper, which observed that “the robots did not like people with dark skin.” Balkin observes that this statement is an instance of the homunculus fallacy—the reporter had incorrectly attributed the racism to the program itself, deflecting the responsibility for the algorithm’s behaviour away from its designers, and their possible biases¹¹. Under such circumstances, a productive debate may have been shut down, since a critique of a piece of software, as opposed to its designers, deflects culpability.

In Barthes’ essay “Death of the Author,” the Author (who is tacitly assumed to form the meaning of the work, and whose identity propels that meaning), is a bit like a homunculus. The homunculus seems to explain the behaviour of the machine by granting it an anthropomorphic spirit, as does Barthes’ Author, in that the Author grounds the text’s meaning. However, neither the homunculus nor the Author actually exist, if Balkin and Barthes are to be believed. Both argue that the source of meaning is the individual, the reader who constructs the text in her mind, and the user who filters the algorithms’ behaviour through the lens of the social, to make

10. The website for the contest is located at <http://beauty.ai>

11. The identification of ethical problems arising from algorithm design is a growing area of research. For example, a report published by the Australian Council of Learned Academies entitled “The effective and ethical development of artificial intelligence: an opportunity to improve our well-being” 2019, is a comprehensive study of the implications of emerging Artificial Intelligence for humans, and of the potential risks associated with the development of such digital systems.

sense of it. If the death of the Author is the birth of the reader, as Barthes says, then what is left behind when the homunculus is dispelled? The machine still acts, but those acts do not come from inside the machine, instead they may come from the agencies acting through it. The death of the homunculus, then, might be thought of as the machine becoming transparent, if we finally can see one another reflected in its functions.

An awareness of human agents' spirits active in the unfolding of a program such as a video game recalls the previously-introduced idea of the composer who hangs like a ghost over the performance of his or her work, as the root of the composition's meaning. This understanding of concert performance is embodied in the idea of the "werktreue." Yet the deceased composer may be excluded from an active role in a performance, where she cannot interact except through the notes she inscribed on the score. Software, in contrast, preserves not just sequences of events but also patterns of action, values expressed through functions that evaluate and make choices in a biased manner, especially in more nonlinear, open forms. A video game designer might ask: "should the player do this certain action, how do I respond to that?" The coders' spirit is reactive, a quality which I feel distinguishes the coder from Barthes' Author. The reader construct their meaning from a static text; but the software user can, under certain circumstances, interact with a program actively, testing their choices through encountering affirmation or resistance that reflects the values of the programmer. If the programmatic homunculus is itself a fiction, what can we learn from our interactions with software? I believe it is the trace of other human agents, whether the coder's, other users', or those of the designers of its framework. The death of the homunculus is not the birth of the user, but the realization that we may act in a social context when we interact with a program—unlike Barthes' dead Author, the coder yet lives in her creation, and under certain circumstances may become an object for the fostering of empathy.

2.6.4 Augmented Reality

In Augmented Reality (AR) applications, computers superimpose digital constructs upon the physical world around us (Azuma 1997, 355). This form can be contrasted with "virtual reality," (VR) in that AR environments blend the physical

and digital world so that the user can see both, where virtual reality create immersive, synthetic environments where the user cannot see the physical world surrounding her. AR, however, need not be restricted to superimposing visual content. For example, “Proprius,” an interactive augmented reality composition by Zaynep Ozcan and Anil Camci (Ozcan and Camci 2018), in which the soundtrack sonifies the behaviour of organisms interacting across various levels of the hierarchical food chain. The score is partially governed by the listener’s location in the installation space, making the listener an active agent in the ecosystem, and creating a feedback loop between composition and listener (5-6). I will discuss this work in more detail in Section 2.8.11, but in terms of the current discussion it points to a merging of algorithmic music and augmented reality techniques and technologies that may provide stimulating new territories that composers and programmers can explore.

While “Proprius” is an auditory augmentation of space, most other AR applications also contain a visual component. I will introduce one of the most famous, “Pokemon Go” in Section 2.8.11, because it serves as an inspiration for my own later doctoral work in which I endeavour to bring the adaptive algorithmic score into the physical world—see Chapter 7 for a description of my locative score “Shards of Memory.” There are many audiovisual augmentations of space available to the public, particularly as mobile phones begin to include sensors such as LiDAR (Wilson 2020). “LiDAR” stands for “Light Detection and Ranging,” and deploys lasers to measure distance and depth. The technology is available on Apple’s 2020 iPad pro, and is expected to be featured on two models of upcoming iPhone 12, allowing these devices to scan their surroundings more accurately and deeply, which in turn increases the creative options for app designers. AR artists, including myself, may find in emerging AR technologies new canvases upon which to create, blurring the real and virtual in a way that I have found introduces a tension into artworks. Video games as well are exploiting the technologies, with one of the first LiDAR-specific apps being Niantic’s “Hot Lava,” an adventure game that has the player dodging rivers of lava in their immediate physical surroundings (Wilson 2020). While it is uncertain whether this AR app will replicate the success of “Pokemon Go,” the first and possibly only AR app to date that has been widely popular, as an artist I look upon AR as an opportunity for the interdisciplinary artist working between programming and composing to establish new types of experience and a new, dynamic musical language that may

break new musical ground. Technologies like AR are important not just as technical achievements but as invitations to extend, subvert, redefine, and/or supplant earlier forms. To me, AR is the nebulous ground where the composer gathers the film score, video game score, and classical concert score, and casts them into the real world, to discover surprises that appear once they are unshackled from the mooring points of written score, concert event, or video game.

2.6.5 The Video Game as a Social Space

Playing video games need not be a solitary experience, in the way that reading a novel might be. A game might simulate an environment and characters within it with whom the player can interact. Those characters can even be, like the player, the in-game manifestation of other people, puppets with whom a player can communicate, interact, and can engage in social dynamics of friendship or control. In a study of player social interactions in games of the MMORPG genre¹², Helena Cole and Mark D. Griffiths observed that such games could form social environments where players could engage in forms of interpersonal relationship commonly observed offline—making friends, experiencing attraction, and sharing sensitive information in conversations with other players (Cole and Griffiths 2007, 581-582). Such research provides evidence of a human capacity to relate to others mediated by virtual constructions located within the digital domain. The quality of such online relationships has been examined. For example, a pioneering examination of social networks was carried out by boyd and Ellison in 2007, where they identified the individual account holder as the focal subject of social networks, novel development compared to more established online gathering places, which were typically focussed on topics of common interest to members rather than the members themselves (Boyd and Ellison 2007, 219). When video games serve as a social space, such interactions can be rewarding, with evidence that playing video games cooperatively can promote healthy social activity, though this is mitigated when competitiveness or an achievement-oriented approach is adopted by players (Halstead and Rolvsjord 2017, 1100). Other research

12. MMORPG stands for “massively multiplayer online roleplaying game,” a genre of video game where players congregate through the internet in virtual worlds where they can meet and adventure with other players.

into the social roles of video games has identified their value as a site of activism, particularly in a subset of “persuasive games,” where players learn a game’s rule system as a means to introduce or reinforce concepts in a player’s mind (Anderson-Barkley and Foglesong 2018, 256-257). The construction of a virtual self may thus be an invitation to grow or otherwise challenge oneself socially.

What do such constructions look like, and how might they form connections between computer users? Aleks Krotoski and Ben Hammersley explore the concept of “identity” and its transformations in the digital medium (Krotoski and Hammersley 2015). They define a “digital identity” as a “set of data that acts as a unique reference to a specific object, where that object can be a person, a thing, a concept, a group, or any other definable entity” (1). The main role of the digital identity, they go on to say, is authentication, to allow a user to determine that an entity is what it is believed to be, a binary condition that can only be true or false (1). Humans acting in virtual domains, however, are associated with constructs that are more complex than a digital identity. The “online identity” is an expression of the self, mediated by a digital environment, which, unlike the digital identity, can be fluid (7). Krotoski and Hammersley argue that this mediation changes the nature of the expression of online identity, causing its form to diverge from the expression of self that an individual projects in the real world (3). They argue that this split between online and offline selves can encourage the formation of a “super-me,” an idealized self that is projected towards the digital audience, and that the user might use to conceal “negative” attributes (4). The construction of the online identity is also restricted in some ways. The construction of the online self might be limited by the range of attributes by which that user can describe herself (5). A profile might require the user to identify their gender, for example, but only give a choice of “male” or “female,” a binary that some people do not observe. Through implementing this filter, the designers of the online environment influence how an individual’s online identity is constructed.

“Avatars” are digital objects that link human players with the gameworld in the form of an interactive, graphical, and social representation of the player (Banks and Bowman 2016, 2). We may identify with the avatar, seeing it as similar to or the same as ourselves, and feel responsible for its wellbeing, where we know that we can control its fate, and are obliged to ensure its prosperity (3). In some genres, such as role playing games, the user may encounter NPCs, non-player characters controlled

by the computer with whom the player can interact through her avatar. NPCs can have their own histories, and exhibit a wide range of reactions to the player. When they join the player as companions to her avatar, they can be integral to helping the player achieve her goals (Coulson et al. 2012, 6). Complex interactions can form between player and NPCs, and this can enrich the experience of the game, by adding the dimensions of real world social relationships, and these relationships can evolve as the player's actions extend the inbuilt histories of the NPC, a history of which the player increasingly shares a part (7). In a study of the RPG "Dragon Age: Origins," Mark Coulson et al. found that the NPC characters who were designed as companions could be liked or loved by players, and that certain characters were more popular than others, which the authors take as evidence that the online relationships were similar to offline ones (14). Other research has found that, for individuals with social phobia, social interaction online worlds can become a substitute for interactions in the real world, and that high identification with an avatar can be an indicator of addiction to internet gaming (Sioni, Burleson, and Bekerian 2017, 14-15).

I see the relevance of this research to algorithmic music as follows. Video game music is not a passive art music that is merely observed. Rather, in Karen Collins' words:

... Game players play an active role in the triggering of sound events in the game (including dialogue, ambient sounds, sound effects, and even musical events). While they are still, in a sense, the receiver of the end sound signal, they are also partly the transmitter of that signal, playing an active role in the triggering and timing of these audio events. (K. Collins 2008, KL.85).

By extending some degree of control over music to the player, the agency of the player becomes active in one of the core engines governing affective qualities of the gameplay experience. Additionally, if social interactions with NPCs in a video games can resemble social interactions in the real world, then there is evidence of an uncanny meaningfulness emerging from human experiences in the virtual domain. Creating a video game can take on aspects of modelling a society, and that society, when the player invests their spirit in it, can influence them as social beings—a feedback loop. Agentic acts in a virtual world, insofar as they affect our avatar, can affect

the player behind it as well. While this phenomenon might support the contention that video game design is a process whereby the designer, as an agent, can exert power over players that extends beyond the gameworld and into the real one, in the case of algorithmic music, it suggests that musicking may not only be possible in a virtual world, but that the impact of musicking upon participants (including players mediated by avatars and the code-animated agency of the game and score designers) may extend to the real world as well.

I explore this application of agency in my work “Shards of Memory” (Chapter 7).

2.7 Relevant Terminology and Concepts

This chapter begins with a quote from Brian Eno, who likens the composer to a gardener who plants seeds to watch them grow. The blooming garden is facilitated, but not controlled, by the gardener, who is seen as content to set initial conditions and then step back to discover the beauty emerging from the potentials of nature. It is an apt framing of the agency of the composer who uses program code to create conditions from which music can emerge outside of her direct control, and finds satisfaction in witnessing the process.

This section introduces key terms in the field of computer music, particularly those that are relevant to the creative work presented in this dissertation.

2.7.1 Media and New Media

Both written notation and coding environments are “media” within which music can be evoked. On the surface, however, they seem to be very different constructs. Understanding their nature as media is helpful to understand what they mean for the composer who acts through them. What is a medium, then? Marianne Van de Boomen provides a helpful, general description:

Let us begin with the most basic assumption about what media do: they mediate. In general discourse... ‘mediation’ pertains to acts of negotiation, intervention, and alignment between distinct parties or domains. Mediation occurs when conflicting or match-seeking parties... call in a

third-party mediator to settle a problem, or to negotiate a match between dating or wedding partners. Media as mediators do the same thing: they connect separate domains—whereby it is assumed that there is a domain of objective reality, separate from the domain of subjective human understanding that has to come to terms with this reality. Media then are conceived as third-party in-betweens, instances that mediate between these two domains. They enable contact with the world by intervening between ourselves and reality. (Van den Boomen 2014, 77)

From the previous quote, it is possible to infer that when the technologies that mediate between us change, the quality of our representations of the content of the objective world may change as well. Of course, now that the digital computer has been introduced into the collection of media upon which artists and the general public may draw, the possibilities and complexities of making and sharing art have shifted. When we approach art made with or through computer technology, we enter the fluctuating ground of new media art. The term “new media art” has been contentious (Quaranta 2013, 23)¹³, because the area is still relatively new and its boundaries remain in flux (23). Edward Shanken’s view of new media art embraces Alan Turing’s idea of the computer being a “universal machine” (Shanken 2016, 472), one that can “emulate the specific functions of any other dedicated device” (472). Shanken sees this property of computers as a signal of their capacity to disrupt medium specificity, affording new media artist an opportunity to engage in “a strategic questioning of the nature of media in artistic, technological, and social contexts” (473). He notes that Duchamp’s readymade “The Fountain” (1917), a urinal intended to be affixed to the wall of a gallery, was rejected by an exhibition of the Society of Independent Artists (474), a rejection that exposed not the limitations of Duchamp’s artistic ability, but the limitations of what might constitute art.

New media was founded in computerization (Santaella 2015, 137). Lev Manovich pinpoints the emergence of new media in the encounter of computation and cinematic film occurring in the development of Konrad Zuse’s (1936) foundational computer system (Manovich 2001, 25). Zuse’s machine, according to Manovich, was the “first working digital computer,” using discarded celluloid film strips as a medium into

13. Quaranta lists terms such as “computer art”, “digital art”, “media art”, “electronic art”, and “cyber art”, to name but a few 2013, 24

which binary codings were punched. In repurposing nondigital cinematic film to the task of holding binary data, “cinema became slave to the computer,” a creation story for the computer in which the offspring consumes its progenitor (Manovich 2001, 25).

In his influential book, “The Language of New Media” (Manovich 2001) Lev Manovich lists five definitive properties of new media objects:

- *Numerical Representation*: New media objects are represented as numbers, and can be formally described in mathematical terms suitable for algorithmic manipulation (27-28).
- *Modularity*: New media objects contain multiple units or modules that together form a functioning whole (30-31).
- *Automation*: New media objects can act independently of the direct, immediate control of a user or programmer (31).
- *Variability*: New media object can manifest in many forms, since their components are represented digitally and may be subject to algorithmic manipulation (37).
- *Transcoding*: This term refers to the translation of information from one format to another. The new media object is layered, incorporating a “cultural layer,” understandable to humans, and a “computer layer,” comprising numerical data and algorithmic processes. Transcoding here is the process by which the two layers influence one another (45-47).

What, then, is the nature of the relationship between the digital object and the real object? What distinguishes a performance on a violin from one coaxed from a sample library in Cubase? Or a melody written with quill and parchment versus one generated algorithmically? An answer may be found in Marshall McLuhan’s famous contention that the content of a given medium is always another medium (McLuhan 2013, 24-25). Chains of media emerge, where successive media, energized by new technologies, encapsulate those of the past. For examples, writing contains speech, print contains writing, and an audiobook contains print. This process in which media enfold other media is known as “remediation,” and it provides a more general perspective on digital media that complements Manovich’s.

Remediation

Jay David Bolter and Richard Grusin refer to the relationship between the new, digital media and the traditional media upon which they draw as “remediation,” which they consider a “defining characteristic” of digital media (Bolter and Grusin 2000, 45). Remediation occurs when one medium represents another medium (45), and where this occurs, a “double logic” is evoked: the cultural desire to “multiply its media,” and the desire to “erase its media” (5). The latter objective is described as “immediacy,” to connect to reality through a medium transparently and without distance (5), as we might see in contemporary live-streaming or in following the minutiae of our favourite celebrity’s social media posts. The other objective is characterized as “hypermediacy,” the multiplication of media seen in, for example, the hypermediated CNN website, which even today juxtaposes text, graphics, video, modular design, and a system of web links that together create a multi-medium presentation of the news (9). Immediacy and hypermedia appear contrary in their aims, but Bolter and Grusin see both as active in the same contexts, a paradox that they describe as the “double logic” of remediation—that it identifies a desire to be both transparent in the representation of the encapsulated medium but also to draw attention to the medium and its possible multiplicity (Bolter and Grusin 2000).

Metamedium

Through the lens of remediation, the computer may seem to be amorphous in its identity, a chameleon of media that is as it is used. Nevertheless, the computer, arguably, can be described in terms of its own identity. One such term, proposed by Alan Kay, describes the computer as a “metamedium” (31 Kay and Goldberg 1977; Barnes 2007, 18). A metamedium, according to Kay, is “active,” in the sense that it responds to its users’ interactions. This, Kay notes, is a capacity that transcends most earlier technologies, since earlier technologies such as film camera or a pencil do not adapt to the wishes of the individual wielding it, beyond those applied through direct physical manipulation (Kay and Goldberg 1977, 31). Manovich builds on this idea, explaining that the pioneering scientists and engineers who were designing early computer technology, such as Kay and Engelbart, were not designing computers as an emulation-focussed tool—they were consciously attempting to create a new medium

possessing new properties that would “allow people to communicate, learn, and create in new ways.” (Manovich 2013, 92) Within the software environment, there is no comparable “analog” medium like a pencil. It was, and is today (on the edges of computer science and software art), a *tabula rasa* that affords its users opportunities to create foundational work substantially distinct from objects created using earlier media.

The Importance of the Medium of Software

Media theorist Lev Manovich believes that software itself is a significant medium, stating that “software has emerged as the main new media form of our time” (Manovich 2014). He expands on this in idea his book “Software Takes Command:”

Software has become our interface to the world, to others, to our memory and our imagination—a universal language through which the world runs. What electricity and the combustion engine were to the early twentieth century, software is to the early twenty-first century. (Manovich 2013, 2)

I appreciate Manovich’s observation that the appearance of software, and its quick spread into ubiquity, has had a profound impact upon the construction and expression of the self. The exploration of software as a medium may be assumed a primary concern of artists who work through code—including in my own works, where the meaning of the works in this dissertation is identified largely in the complex relationship between agents, human and computational, as they encounter one another through the software medium. My doctoral work is wholly dependent on the nature of the digital computer as a metamedium, in that it forms a site where actions affecting processes, sound, and imagery, carried out over time and space, work in a unified manner. For example, the syntactic and procedural conventions I deploy to form the code that animates the visually-presented, orb-like avatars in “Shards of Memory” are the same or similar conventions I use to achieve musical and other sonic effects. Yet my ability to sculpt these forms of app content (graphical and sonic) also depends on my experience with music and graphic art that has roots in my practice that began to form long before I made my first digital painting or score. I can compose through code because it remediates music and art, and my earliest improvisations on

my electric keyboard and pencil drawings I fumbled through as a child form tangible roots for what I ultimately produce through code. The nature of the digital computer as a medium is one that draws together creative action, between people and within the diverse expressive skills of an individual, and so becomes an extender of agentic action, social and internal.

When I am asked what instruments I play, and say the computer is my instrument, it is perhaps a diminution of the computer and its role in my work. For me, the computer is the focal point of my creative action, and the prism that reveals the colours of my self-expression. The digital medium is the impetus for evolving my interdisciplinary methods (which are discussed in more detail in Chapter 3), and it facilitates a convergence of all manners of my creative expression, in that it remediates all of the art-making that has led up to my current doctoral works. On this basis, I maintain that composing through the medium of code is not a divergence from notation-based score-writing, but an instance of remediation. All of the coding projects presented in this thesis draw upon my knowledge of and experience with composing music, and my agency, projected as a coding composer, is rooted in the same skills and expression that I apply when acting as a agent writing a score. Thus I consider the computer to be no mere instrument, but rather a layered domain offering opportunities for agentic interplay that can reward the artists who dare to strip away its surfaces in a search for novel opportunities for expression.

2.7.2 Real-time Composition

Carlos Guedes uses the term “real-time composition” to describe a “compositional practice utilizing interactive music systems in which generative algorithms with a non-deterministic behaviour are manipulated by a user during performance” (Guedes 2017b, 162). Guedes highlights certain aspects of real-time composition: it is a performative practice that is interactive and its form is not pre-determined.

The concept resonates with foundational ideas of Joel Chadabe, who, beginning in 1967, developed methods that he called “interactive composing,” through which he introduced “performable, real-time computer music systems in composing and performing music” (Chadabe 1984, 22). Interactive composing, in Chadabe’s model, is a two-stage process: first, a system is developed, and second, it is deployed by a user

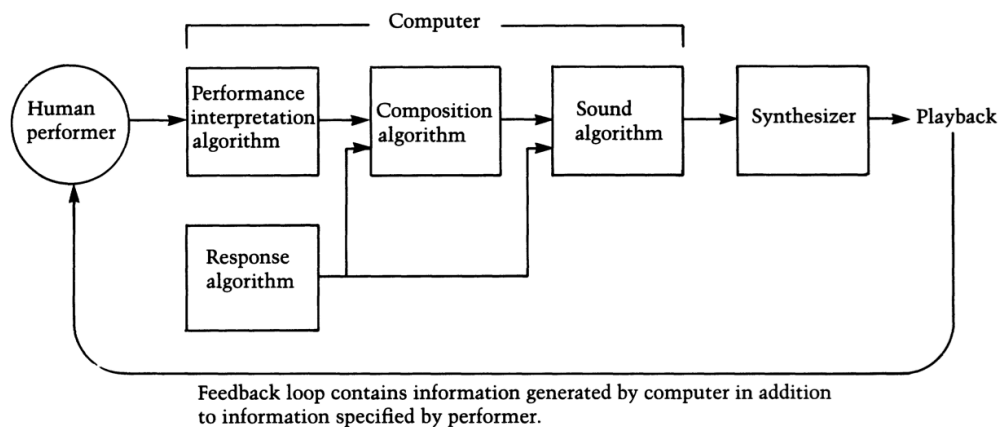


Figure 2.1: Joel Chadabe's depiction of a prototypical interactive composing system (Chadabe 1984, 24).

to simultaneously compose and perform music through interactions with the system (Chadabe 1984, 23). Chadabe presents a diagram of such a work, reproduced in Figure 2.1. It is structured as a feedback loop in which the human becomes a component of a composition-generating system, essentially working as a module alongside or in sequence with a computer and synthesizer. Chadabe’s flowchart suggests a differentiation of roles for the various components of the system. The human here is the performer, and the system’s response to the performer’s actions must be “interesting and informative... [and so] must contain new, unexpected information” (25).

Chadabe emphasizes that the computer’s role is to respond like it is engaged in a “conversation,” in that the computer’s responses to the user’s actions are “recognizably related to the performer’s actions,” yet are also unpredictable (Chadabe 1984, 25). The result is the performer exerts agency over the music that emerges from the system, but not total control. The computer acts as an interpreter for the performer’s actions, and controls aspects of music that are outside of the control of the performer (23), with the synthesizer’s output reflecting a blending of the agencies of the program and its user/performer.

Chadabe says “music and performance are inseparable,” and thus composers and performers are offered a new way of engaging in acts of musical creativity (26-27). In his own work there is a clear blurring of boundaries between the roles of composer and performer—see Section 2.8.3 for a depiction of “Solo” and “Rhythms,” two of Chadabe’s original works which exemplify interactive composing.

Not all automated composition systems may rightly designated as instances of interactive composition. Consider the “Illiac Suite” (1957), the computer-generated composition designed by L.A. Hiller and L.M. Isaacson discussed above. The suite was the first computer-based composing system to use random processes in the composition of a fixed score (Sandred, Laurson, and Kuuskankare 2009, 1). The program used probability tables to randomize the selection of pitches which were then applied within a fixed eighth-note ostinato pattern. Tables specifying probability distributions of melodic intervals were changed every two measures. In terms of the distribution of agencies that combined to craft the suite, the work is not actually an instance of interactive composition or RTC. While the programmers did develop an algorithm that specified musical events, and those events were not directly predictable to the programmers due to the implementation of stochastic processes, the system

constrains the user to the role of initiator rather than interactor, since there is no input available at runtime to engage with the probability tables. Thus the algorithm is non-performable, and as such it is not an instance of interactive composition.

2.7.3 Algorithmic Music

In their introduction to the “Oxford Handbook of Algorithmic Music,” Alex McLean and Roger T. Dean define an “algorithm” as simply “a finite sequence or structure of instructions” (McLean and Dean 2018, 2). The concept of the algorithm was conceived long before the invention of computers, and were applied in music-generation applications such as eighteenth-century musical dice games (also known as *ars combinatoria*), in which compositions were assembled by deciding the order of measures by throwing dice (N. Collins 2018, 68; Jeongwon and Song 2002, 264). Other procedural approaches to include Marcel Duchamp’s “Erratum Musical” (1913), in which notes were drawn from a hat, and Hayes’ technique of splattering paint onto score-paper to determine notes (Cormier 1975, 285). What unites all of these early implementations of algorithms in music is that they can be seen to relocate agency over music structure from the domain of the composer into algorithmic processes.

In the twentieth century, some composers chose to relinquish some control of music structures that were more typically governed by composers’ choices. Yet the incorporation of algorithms into processes of composition do not necessarily negate the composer’s agency. Consider Cage’s work “Music of Changes” (1951), chance processes determined all discernible characteristics of the music (Jensen 2009, 97). Tempo changes, for example, were varied every few measures according to chance operations (Cage, Kirby, and Schechner 1965, 68). This yielding of the specification of musical events does not mean that Cage yielded his agency entirely to chance. He states “Most people who believe that I’m interested in chance don’t realize that I use chance as a discipline. They think I use it... as a way of giving up making choices. But my choices consist in choosing what questions to ask” (Kostelanetz 2003, 16). In his analysis of the work, Marc Jensen refers to this shift in the nature of composers’ agencies: “Far from an abdication of composerly control, the ‘Music of Changes’ is rigorously controlled in every aspect other than its linear organization” (Jensen 2009, 98). The algorithm did not silence Cage, but rather invited him to exert his agency

over a new musical domain.

2.7.4 Stochastic Processes

Iannis Xenakis is a central figure in the integration of algorithmic processes into music production (McLean and Dean 2018, 9). In his 1953 composition “Pithoprakta,” which translates as “actions by probabilities” (Kay 1967, 21), Xenakis specifies musical structures using manually-calculated probability distributions (Luque 2009, 77). Xenakis named this music “stochastic music” (77). This term derives from statistics. A “stochastic process,” according to Kevin Jones, is a “collection of random-variable quantities distributed in space or time,” and “stochastic structure” is a “framework for analysis of stochastic processes” (Jones 1981, 45). In musical composition, stochastic processes are often used to generate, rather than analyze, mathematical variables in a composition. A stochastic structure is associated with an “event space,” a set of possible events that may occur. The selection of these events is governed by a probability distribution (46). For example, the notes of a C major scale might form the event space, and a probability distribution weighting the notes of the C major triad more highly than the other four pitch classes might be used in an attempt to evoke the sense of music written in the key of C major.

Xenakis’ use of stochastic techniques was originally implemented through manual calculations conducted by the composer himself; in 1962, however, Xenakis began to use computers as a means of accelerating the calculation processes within which his stochastic compositional methods were applied (Luque 2009, 77). Indeed music composition itself can be viewed as a stochastic process (Liu and Ting 2016, 2) in which the act of composition involves specifying event values (e.g. the pitch of the next note) that can vary. With the computer dramatically speeding the rate of calculations, Xenakis proposed applying stochastic methods in all levels of compositional processes (Serra 1993, 235). To this end, Xenakis created the program GENDY3, in which all aspects of its compositions, from music structure to the output timbre, were governed by automated, stochastic processes (237). Stochastic works such as GENDY3 automate and accelerate aspects of composition (237), challenging the traditional role of the composer by shifting the locus of their agency from direct specification of note events to the specification of probabilistic processes which synthesize

the composition.

As an algorithmic composer who was initially reluctant to give up control over variables afforded by the form and grammars of music notation (even as implemented in digital composing environments), the idea of building generative process, rather than directly specifying notes, raised in me a certain discomfort. Xenakis' use of stochastic techniques is instructive in this regard, as his work demonstrates that chance does not necessarily obliterate choice. Indeed, in my own work "Shards of Memory," the user triggers GPS-bound events in a manner that is non-deterministic, essentially making the unpredictable path which the user takes as an input to a stochastic process generating the score. My role as a programmer and composer becomes in part to anticipate all such paths, and sculpt the underlying algorithms in such a way that the path of any user will function towards my general aims; however, it is the user's mobility that sets the work in motion, and in that regard, the expression of the piece does not exist without their engagement in triggering the various stochastically-governed musical events that populate the augmented landscape.

2.7.5 Generative, Reactive, and Interactive Music

In an article considering potential avenues for mass-market distribution of algorithmic music, Yuli Levtoy defines three types of algorithmic work (Levtov 2018, 628).

Generative music. In generative music, no factor external to the software producing the work can influence its production. That means that the listener becomes a passive receptor of the music being produced, and has no agency over its form. Generative systems are not necessarily deterministic, and can incorporate internal sources of variation such as randomization or initialization with user-provided datasets. Such sources of variability affect the unfolding of their algorithms, but this variation occurs within the internal actions of the algorithm, and cannot be changed after processing begins (628).

Reactive music. This type of algorithmic music is very similar to generative music, but is distinguished by the fact that its parameters can be influenced by factors outside of the algorithm, such as the environment of the listener. Examples might be real-time sensor streams providing measures walking pace or heart rate, or local weather

information sourced from the internet. Crucially, in the case of reactive music, the listener remains passive, exerting no control over the unfolding music that could be ascribed to intent (Levtov 2018, 628-629).

Interactive music. In this form of algorithmic music, listeners actively exert their agency over the unfolding music, working through an interface to deliberately play a role in directing the compositional processes carried out by the algorithm (629). Here, the listener crosses over the creator/audience divide, with the agency of the listener joining that of the designers to govern the unfolding of the work.

In my doctoral endeavours, there are examples of all three approaches. Q-Chords (see Chapter 5) is a purely generative work, for once it is configured and set in motion, no interaction with the system on the part of the user is possible—the resulting chords are subject to chance operations, but this is outside of the immediate control of the user. In “Chief Paskwa’s Pictograph” (Chapter 6) visitors to the exhibit influence the unfolding of the music through shifting their location within the exhibit space, but this control is intended to act invisibly, and not to invite their direct engagement with the motion sensors to project their intent. In “Shards of Memory” (Chapter 7), however, the user exists within the augmented environment layered over the physical world in the form of a melodic avatar that follows their movements. This avatar is an interface that shadows in the world of the installation the user in the real world. This is the source of the user’s contribution to the score of “Shards of Memory,” and the feature that makes the work interactive. Since interactivity is an important idea in this work, I will delve a bit deeper into the concept in the next section.

2.7.6 Interactivity

Many computer music forms have been described as “interactive,” but the definition of this word is open to debate. In his 2001 treatise on New Media, Lev Manovich goes as far as to say that “the concept is too broad to be useful” (Manovich 2001, 55). I feel that the term can still be useful, if its definition is clarified and limited in scope. One useful definition is provided by Tim van Geelen, who characterizes interactivity as involving “a message with multiple relations to previous messages and to relationships between them” (Van Geelen 2014, 561). This definition can be used to

distinguish interactivity from reactivity. Reactivity is exemplified in the act of turning a volume knob, where there is only one relation to a message: the knob turns, the volume changes (Van Geelen 2014, 561) Multitouch on an iPad may also be reactive, even if each point of contact conveys a different message, in cases where there are no influences between messages themselves (such as when one plays a virtual piano keyboard). Strumming a guitar is also a reactive act, since the performer can vary only one or two dimensions of the sound. The choice of what to play, however, can be interactive, where the choice is based in part on memories of previous messages, specifically the content of the performance leading up to the current moment (562).

David Rokeby provides another view, relating more to subjective experience than technical criteria. He states that “technology is interactive to the degree to which it reflects the consequences of our actions or decisions back to us” (Rokeby 1995, 133). A volume control knob would be minimally interactive, since it only reflects the choice of dynamic level to the person who turns it. He quotes Itsuo Sakane, a Japanese journalist and curator, who states that “all arts can be called interactive in a deep sense if we consider viewing and interpreting a work of art as a kind of participation,” a view, Rokeby notes, echos Marcel Duchamp’s’ declaration that “the spectator makes the picture” (134). Through these lenses, the artist acknowledges the agency of the audience, and at some point in the development of the artwork, chooses not to choose from whatever possibilities remain in the formation of the work, yielding control of such variables to “a complex, indeterminate, yet sentient, element: the spectator” (137).

Rokeby provides descriptions of possible models of interaction between an artwork and interactor. *Navigable structures* are spaces designed by an artist for the spectator to explore, granting the spectator a sense of freedom which the artist might constrain (138-141). An example might resemble a labyrinth, or an augmentation of physical space. *Invention of Media* occurs where the artist designs means for the interactor to create objects within the work, but where the form of these objects is constrained by the work (144-145). In *Transforming Mirrors*, the “idea of the mirror is explicitly invoked” (145), and a representation of the interactor (such as a silhouette, mirror image, or shadow) is manifested by the computer controlling the artwork. The tension formed as a result of divergence between the interactor’s self-awareness and the form of the behaviours of the mirrored self can challenge the interactors representation

of herself (Rokeby 1995, 146-147). *Automata* are artificial entities, programmed to behave in ways that explore the forms and limits of human behaviour (151). This latter form is less concerned with mirroring the spectator, and might use interaction to reveal qualities of human social environments (151).

Interactivity is a feature of a system that can allow agents to gain control of it, and even encounter one another through it. A satisfying implementation of interactivity is an objective of my work, since this thesis is concerned chiefly with the impact of coding on the interplay of agencies in musical settings—particularly as it applies to the composer. I had to divest myself of the trappings of a venerable composer, particularly the status of sole owner and author of my scores, and instead embrace the possibilities of agentic interplay afforded by interactive inputs into the control systems governing the music. In this sense, a quest for a satisfying implementation of interactivity is at the root of all of my creative work completed for my doctoral program. If interactive technology reflects us back to ourselves, I had to learn the value of bringing forth a mirror for my audience, which, in the end, helped me to understand my own motivations and biases more clearly.

2.7.7 Nonlinearity in Music

Conceptualizing music as a nonlinear process rather than a linear one is a complication associated with algorithmic music, so a brief discussion of these terms is warranted.

Linearity is a property of sequences whose elements occur in a defined order. (Kaae 2008) This form applies to much traditional Western music—we can anticipate each note of a performance of Beethoven’s *Moonlight Sonata*, for example, since these notes are set down in the score of the work. Such traditional scores define events along a linear timeline (Magnusson 2014, 268). Most digital audio workstations that feature notational or piano roll facsimiles of a notated score, including Cubase, encourage a deterministic, linear approach to composition, despite being implemented as coded programs running on a computer (269). Yet musical notation itself can be viewed as a code, a “systematic format of instructions for the performance of musical events,” according to Thor Magnusson 2014, 269. When this language is presented on a sheet of paper, it may often, but not always, be seen as indication that the

events represented can be repeated on some level, even in different performances by different musicians (Magnusson 2014, 270). When notational conventions are digitized, however, a challenge to linearity emerges.

A nonlinear system grants time a dynamic quality. Kaae defined dynamic music as “music which is in some way composed or put together in real time by the computer” (Kaae 2008, 1802). The key idea here is that the composition of the music is not pre-determined before the performance, as occurs in much Western tonal music, for example. Non-linear music can deploy a process of variation that takes time to unfold, as opposed to an ordered sequence of events that have predictable relationships with one another because they are set down in a score (1816). On the relationship between linear and nonlinear music, Myron Kruger, an early innovator of virtual reality systems, states:

“a composer [of traditional linear music] who typically defines a detailed sequence of events [is now] composing a sequence of possibilities, many of which will not be realized for any given participant who fails to take the particular path along which they lie.” (Krueger 1977, 431)

In broad technological terms, music becomes dynamic in two ways (Kaae 2008, 1976). The first is through variability. This involves selective processes specifying events, such as through implementations of probabilistic stochastic methods. The second is adaptability, where the music reacts to stimuli of which it is aware that are external to the code governing music structure. Both of these techniques have long been utilized in video game scores,

Despite the existence of many possible forms of a nonlinear work, the experience of interactive artworks can appear to be linear to their audiences, since they necessarily observe the unfolding of music over linear time, where musical events may seem to fall into an inevitable timeline (1855). For the designer of an interactive music system, however, it is beneficial to move beyond this viewpoint, thinking of the interactive artwork as a series of potentialities, which Jesper Kaae describes as “like seeing time from the inside” (Kaae 2008). Indeed, nonlinearity used in video game contexts has been described by Richard Rouse as the attribute that gives interactivity meaning, for, without it, the video game effectively becomes a film (K. Collins 2008, 96).

Before my doctoral program, I conceived of composition as a process aiming to craft a linear sequence of musical events, to create an authoritative document of the progression of my thoughts. Kaae’s description of nonlinear music as granting an additional dimension of time to music is elegant and points to the value of transitioning music from linear models to non-linear. I can only imagine the shock if a population of beings existing in a two dimensional plane were suddenly introduced to a third dimension—to go over one another, rather than around. The world would expand, models of space and possible behaviours within it would become limited and need complete reconsideration. That is what nonlinearity brings to music: the ability to go over instead of around, where, if a path is not possible in this instance, perhaps it will be in the next. The printed, notated score is capable of this in a hollow form, by not specifying every event exactly. One performer might read a line as loud, the next not quite as much. In that ambiguity lies the expression of the score, that the composer could not specify. And, as ambiguity in a score grants expression to the performer, the introduction of an additional dimension of time in the algorithmic score may bring forth an even more dramatic ambiguity, one that may admit a multitude of agencies: listener, performer, user, programmer. In algorithmic music nonlinearity is the beginning of a vast uncertainty, and so becomes the means through which outside agents can begin to be heard in the music.

2.7.8 Adaptive Music

Another term associated with interactivity in sound is “adaptive music.” Commonly applied to dynamic video game scores, adaptive music can be distinguished from interactive music by its action in relation to the player. Interactive sound responds directly to the player, and the player is aware of their impact on the sound, whereas adaptive sound changes in form discreetly to enhance gameplay, responding to player behaviours without drawing the player’s attention to those changes (Clarke 2007). In a sense, the distinction between adaptive and interactive music parallels the distinction between a film score, which may be intended to function in the viewer’s unconscious, and the concert work, upon which a listener is expected to focus when it is performed.

Karen Collins distinguishes interactive from adaptive audio forms in a slightly

different manner, explaining that interactive sound should be defined as a audio response to the direct input by the player (e.g. sound effects like footsteps, gunshots triggered by the player, and the distinctive sound of the Nintendo character Mario jumping), whereas adaptive audio responds to the game state (e.g. changes to emotional expression tied to player health or location, and the speeding up of musical tempo as the time to finishing a level is reaching zero) (K. Collins 2008).

In the context of an algorithmic musical work, the dichotomy between interactive and dynamic sound raises questions for the composer. While a musical work need not be associated with content in other modalities (such as the visual component of a film or the movement of a dancer), I find that, when I compose, I do not aim to have every component of the score foregrounded at once. My scores, like many other composers', are layered with material that ebbs and flows in terms of whether it is the focus of the arrangement at a given instant. In my algorithmic music, this idea of foregrounding that I apply in my linear compositional practice merges into questions of interactivity versus adaptivity in music. In "Chief Paskwa's Pictograph, for example, the sound world was intended to be largely adaptive, in that I desired the music to support the experience of the exhibit, rather than overshadow it. I did not want the patrons moving through the exhibit to be overtly aware that they could trigger musical elements by moving around. Were that to happen, it might seem like hypothetical video game that drew players out of its narrative by allowing them to become aware that they could "perform" the score through their in-game actions. If this disrupts the narrative, or the sense of immersion, then the score is not functioning as it should, I feel. So, "Chief Paskwa's Pictograph" was almost entirely an adaptive score. With "Shards of Memory," however, I wanted the users to find themselves in the game world and its music. To this end, the users had very clear control over their musical representation in the augmented world, literally chasing a melody to represent them and interact with objects in their behalf. This is interactivity. There is adaptive music in the app as well, however, as there is continual underscore designed to support the visuals, foregrounded musical elements, and player behaviour. This score is adaptive, and functions in response to the user in subtle ways. In its combination of adaptive and interactive audio, "Shards of Memory" is the closest for my works to a video game score, with the important caveat that the music in the app is foregrounded, rather than supportive of other visual, narrative, or social elements.

2.7.9 Comprovisation

Trevor Wishart describes his work in electroacoustic studio as a “slow improvisation,” a concept by which he connects the behaviour of a composer to that of an improviser (Vassilandonakis and Wishart 2009, 10). He has described his experience creating, or composing, in the analog electronic music studio as a “slow improvisation” (10). Wishart’s personal creative process typically occurred over three stages: a process of collecting sounds, a period of transforming those gathered sounds, and then a definition of the structure of the piece (13). Improvisation, says Wishart, is simply a part of the compositional process, located for him mainly at the material-generation stage where audio was sourced and manipulated. (10) By working with the material in an exploratory manner that tended to lead to unpredictable outcomes, and by using such outcomes as a stimulus for growing his library of sonic material, Wishart integrated composition and improvisation in his own methodology.

Wishart joins other scholars in the field of experimental music in regarding composition and improvisation as connected, rather than as incompatible poles on a continuum of planning. Sandeep Bhagwati (Bhagwati 2008) and Richard Dudas (Dudas 2010) explore the relationship between composition and improvisation as identified by the term “comprovisation”. Observing that Wishart rightly identified the improvisational aspect of studio work, Dudas characterizes one form of comprovisation as a play-like process of generating material to be used in subsequent compositional processes (30). The other comprovisational method is that of composing an instrument. He defines an instrument after Atau Tanaka: “a self-contained and autonomous sound-producing object that enables a musician to perform in a live situation” (Tanaka 2009, 30). Computer music performance systems, while they may fit this definition of an “instrument,” can extend it by incorporating a core compositional structure. This compositional structure is both the score of the computer music work while also defining its use and behaviour as an instrument, hence the term “comprovisation.” Bhagwati provides another view, positioning comprovisation as a model of musical creativity reflecting equilibria situated between total composerly control and “totally free-floating forms” (Bhagwati 2008, 4). In his creative practice, he explores forms of action between written music and comprovisational approaches in his own practice,

for example by giving performers options that allow decisions to emerge from the conditions of the performance, for example, from the interaction between musicians or with reference to an independent process (Bhagwati 2008, 5-6). Comprovisation here is understood to involve a “swift ongoing mediation between notation, improvisation and intervention by peers and conductor(s),” where the musician is an interpreter who, unlike in previous art music practices, the interpretation is no longer focused on the composer’s imagination (6).

Dudas observes that improvisation has been “much-maligned” in the Eurocentric classical music sphere which he relates to the predominance of recording technologies that create a fixed copy of events which were not fixed in their unfolding (Dudas 2010, 2). He links these ideas to Andrew Brown’s argument that such recording technology “frozen sound,” and, by virtue of its prominent role in facilitating and disseminating early electroacoustic music, the concept of fixed scores had become entrenched in the community of computer music creators (Brown 2002, 29). This perspective reinforced the assumed separation of composition and improvisation, despite the fact that improvisation was a commonly practiced skill set of musicians and composers from the Middle Ages through the Baroque Era and was even widely and formally practiced during the early years of electroacoustic activity and to this day in the form of jazz. (Dudas 2010, 29)

Dudas quotes the Russian composer Igor Stravinsky’s colourful description of composers who would “often forage aimlessly like animals in order to seek out new musical territory” (30). Comprovisation may be observed in such foraging behaviours, where the artist improvises to generate a well of ideas before recording their work in score or audio format. In an analysis of his own improvisational computer music practice, Joshua Mailman explains that improvisation has been an activity located between composition and performance, drawing substance from each. In turn, Comprovisation, Mailman suggests, “straddles the divide between composition and improvisation.” (Mailman 2013, 351) In the context of coding generative music systems, Mailman positions the method as a hybrid of composition and improvisation. From composition, the improvising new media artist draws a tendency to “compose music-generating algorithms as guided by aesthetic concerns,” and the inclusion of a “planned choreography of physical movements.” From improvisation, it draws

three other aspects: its potential for the creator to use “spontaneously decided physical movements,” the manner in which “planned (choreographed) movements may be spontaneously ornamented with expressive nuanced deviations,” and that, in execution, a “quasi-stochastic algorithm may be regarded as ‘improvising’ since its determining of certain details cannot be predicted in advance.” (Mailman 2013, 351-352) In other words, improvisation as carried out in the design of interactive artworks unites planning with performance, predictability with spontaneity, and rewards the artist who engages in their art-making processes with a spirit of experimentation and an embrace of unexpected outcomes with the possibility of finding inspiration in chance events.

I consider my own creative process to use forms of improvisation, applied both in my compositional and in my algorithmic creative work. Indeed, improvisation is the over-arching method that links my creative practices in the multiple media where I work. As the computer is almost always my interface with the works I produce (including concert music, soundtracks, algorithmic music, digital art and writing), my means of creating in each of these forms depends on my ability to quickly sketch an idea, edit, cut, paste, save, and delete it at will. I find these affordances that the computer offers to sustain the equivalent of a “slow improvisation” where my strength is in making the little connections between ideas that gradually form works of which I generally could not predict form at their inception. By situating improvisation in alignment with an improvisation and performance, I also find an entry point for other agencies into composing and coding processes. After all, a performance of any number of performers greater than one should introduce an interplay between those performers. Likewise, when I compose as a improvisation, I am providing an entry point for other agencies—the designers of my digital audio workstation, the coders of the plugins I use, and the musicians whose sampled performances become part of the expressive tapestry of my score. And, in the case of my coding projects, a improvisational perspective invites me to experiment with configurations that leave structural elements of the music unspecified, to be completed in response to the actions of other agents, such as the user in real time or the contributor of content that is predetermined when the app is run. The interplay of agencies that an improvisational framing of composition and coding introduces is foundational to my practice, and I feel that it draws me closer to being able to claim that the computer is my instrument

(upon which I actually perform), rather than a tool that I wield in a manner that prioritizes my desires, and transmits none of its own, or its designers.

2.7.10 Soundscapes, Soundwalks, and Audio Walks

Over the course of my doctoral program, I have moved towards developing interactive music. In my final doctoral work, “Shards of Memory,” users’ interactions with the content of the work involves the user walking through a physical space, where their mobile phones’ GPS updates govern aspects of an evolving score. In this sense, I draw on the ideas of soundscapes, soundwalks, and audio walks, so an introduction these forms is relevant here.

“Soundscape” is a term associated with pioneering work by Canadian composer R. Murray Schafer. He used it to describe, in analogy to the idea of the “landscape,” the collection of sounds encountered in a physical space (Polli 2012, 258). To Schafer, the soundscape should not be dislocated from its landscape, for soundscapes are as inextricably tied to their location as the objects of a landscape itself, and Schafer described the disconnection of sound from its source as “schizophonia,” (258) manifesting as a sense of “nervousness and aberration” (Thulin 2014, 50). Schafer identifies three components of a soundscape. “Keynotes” are background sounds that, in analogy with the key of a tonal composition, form the base of the soundscape (Thulin 2014). “Sound signals” are foregrounded sounds that carry information and attract attention (Polli 2012, 257). “Soundmarks” are also foregrounded, but, like a landmark, their significance is recognized by the people who encounter the soundmark, such as in the case of a geyser or waterfall (Revoll 2018, 6). Listening to a soundscape may be regarded in somewhat different terms from listening to music, although this is not to say that they cannot overlap. Listening to traditional music, such as a pop song by the Beatles or a Beethoven symphony, is often considered an intentional act, in that the music is foregrounded. Soundscapes, on the other hand, are always present, and may be attended to unconsciously (Polli 2012, 260). That is not to say that soundscapes are not noticed and music cannot be meaningful if listened to passively. Ambient music, for example, is intended to function at a wide range of degrees of attention (see Section 2.8.2 for a discussion of the form).

“Soundwalking” is a practice related to the soundscape, initiated by composer

Hildegard Westerkamp (Polli 2012, 258). Westerkamp defined soundwalking as “any excursion whose main purpose is listening to the environment” (Westerkamp 1974). The main purpose is to connect with the soundscape through engaging in focused listening to an environment while moving through it (Thulin 2014, 136). Such a process is consistent with the perspective of soundscapes, since in the soundwalk the natural form of the sounds are not obscured or separated from the landscape. Soundwalking recalls elements of Pauline Oliveros’ practice of “Deep Listening,” which she describes as “going below the surface of what is heard, expanding to the whole field of sound while finding focus” (Oliveros 1999, 5). In soundwalks and Deep Listening alike, attention to sound reveals insight, and distraction disturbs this process.

Janet Cardiff’s innovation of “audio walks” diverged from the soundwalk by adding an artificial layer to the experience. Cardiff’s first audio walk, entitled “Forest Walk” (1991), required the user to play a taped recording through a Walkman¹⁴, listening through the headphones while traversing a real space (Thulin 2014, 136-137). Instead of listening to and learning from the environmental sounds emanating from the surrounding space, the listener hears a narrative transmitted via the headphones, and walks a predefined path as instructed in the narration (137).

I see a similarity between Cardiff’s audio walks and medium of augmented reality. In fact, Cardiff is essentially creating an audio augmentation of a location, using analog media instead of digital. The distinction between the soundwalk and the audio walk can also be framed through analogy with augmented reality. The soundwalk focusses on the natural, where the audio walk on an interplay between place and augmentation. Augmented reality applications can be designed with either approach in mind—to create a new world in a blending of natural and artificial, or to attempt to transparently reveal aspects of the surrounding real world through augmentations that do not obscure that world’s meaning. Schafer saw soundscapes as under threat by urbanization and other human activities, an indication of the encroachment of human machines into the natural environment in the form of noise (Polli 2012, 257). A similar question may be posed with respect to augmented reality: do its additions to the landscape actually subtract from the landscape? In the case of my own augmented reality app “Shards of Memory,” my goal is to situate the user more richly in an

14. The Walkman was a cassette-tape based, portable music system that anticipated the design of more recent mobile music systems like the iPod or the smartphone

otherwise familiar place by revealing events that were meaningful to other people. In this sense, “Shards of Memory” resembles an audio walk, conceived with interactive and adaptive music at the fore. It is not intended to disrupt the landscape, but rather to reveal the trace of its otherwise invisible social meanings.

The question of how such environments might add or subtract from the landscape itself is one frequently considered by artists of locative media (including myself), so a discussion of such works is presented in the next section.

2.7.11 Mobile and Locative Media

I chose to create a GPS-aware smartphone app for my final project. A smartphone is a computer, of course, and as such devices’ processing power has increased, particularly in the last few years, they have become identified as an emerging platform for algorithmic music (Plans 2017, 51). Some such applications are conceived as “album apps,” works that create an evolving sound field, which may or may not respond to user interaction (51-52). Examples include the aforementioned “Air” app, modelled after the ambient music of Brian Eno, which was preceded by BlueBrain’s “The National Mall” and Bjork’s “Biophilia,” both released in 2011 (Sa Dias 2014, 25). Bjork’s app is distinguished by its lack of linear structure. The user encounters a three-dimensional graphical universe that can be traversed by touching the screen, revealing songs that manifest as distinct interactive audiovisual spaces built around Bjork’s music. “The National Mall,” on the other hand, also responds to user input, but interaction in this app is centred around the GPS sensor (26). The app was active around the National Mall Park in Washington, DC, and users of the app would find that as they walked through the park, their travels would impact the music, changing instrumentation or adding sound effects that referred to objects in the space (for example, approaching the lake might change the piano theme to a harp theme, and nearing the merry-go-round would evoke the sound effects of horses) (26).

Frauke Behrendt 2012 provides an analysis of the use of sound in mobile media from early in its history. He observes that, prior to 2012, locative media had been biased towards “visual, textual, and often map-based interactions,” and argues for a more multi-sensory approach focussed on sound. By so doing, the focus can shift from the device to situated activities (Behrendt 2012, 283). In a previous study (Behrendt

2010, 48-80), he provides a taxonomy of mobile sound art, identifying four forms:

- *Placed Sound*. Artists distribute sound in space and users form their own version or “remix” of the piece through their choices of where to go (Behrendt 2010, 49-50).
- *Sound Platforms*. A platform is provided through which its users can contribute and situate sounds in a space. (57).
- *Sonified Mobility*. These works use user mobility, which might be walking or any other movement, to initiate and influence the output sound (66).
- *Musical Instruments*. The smartphone is imagined as an interactive musical instrument, to be performed by the user through “mis-using” the sensors in unconventional ways like blowing into the microphone or shaking the phone.

Behrendt indicates that these categories overlap, both with each other and outside the realm of sound art (Behrendt 2015). I feel that, in the overlap between these categories, there is a potential to explore complications of musical agencies. For example, “The National Mall” is an instance of placed sound, as it is a GPS-aware app, but it also conforms to the definition of sniffed mobility, in that the user’s movements through space provide the means of discovering the work. The agency of the users is necessary to give the work form, and the linear form of any one instance of the work is partially governed by the user’s agency. If we were to extend the app by allowing the user to contribute or perform musical content, the blurring of the user’s role with that of composer and performer would create an even more complex interplay of agencies.

The most relevant of these four categories to this dissertation is the Placed Sound category, particularly the “Historic” variant, which Behrendt describes as “using locative audio to attach sound, music and narrative with content relating to the past, to locations of historical significance” (10). Behrendt identifies placed sound forms as a variant of augmented reality, forming a layer of sound media (as opposed to visual media, although the two can certainly co-exist) (13-14).

The emerging field of “sonic interaction design” provides a perspective for artists engaging in mobile sound art (14). Sonic interaction design refers to the “practice and

inquiry into any of various roles that sound may play in the interaction loop between users and artifacts, services, or environments” (Rocchesso et al. 2008, 3969). The field sees sound as a medium situated within the “human-artifact loop” (Rocchesso, Delle Monache, and Barrass 2019, 153), a sonic link between the human and the object that I believe can be a site of encounters between agencies.

Locative audio draws upon the idea of the “audiowalk,” introduced above, featuring linear recordings intended to be experienced within a specific location (Hazzard et al. 2017, 1-2). It also relates to the NIME community (“New Interfaces for Musical Expression”), where the phone is conceived as a mobile interface for making music.

My own app “Shards Of Memory” (presented in Chapter 7) is an example of a locative audio work. I feel that it contributes to the field in three ways:

- I establish in “Shards of Memory” an interplay of agencies, mediated by a musical score, that positions locative media as a social site where people engage with the app to engage with one another. This motivation positions the app, and locative media generally, as social experience, promoting empathy, and establishing the medium as a site for musicking.
- I apply my interdisciplinary skills and experience as both composer and programmer to the crafting and coding of the score, which I would argue represents a relatively rare confluence of artistic and scholarly aims. My background as a composer have helped me to helps me to establish a sound world grounded in a rich experience of studying and composing soundtracks, and my programming skills, I would argue, allow me to act as a composer through acts of coding. As I apply this interdisciplinary approach to locative media, I aim to establish a practice of “scoring the world,” beginning with the University of Regina campus.
- I apply locative media technologies and approaches towards forming an archive, in which events are recorded and presented in the place that they occurred. This innovation allows my AR app to overlay history over space, which, in future versions of the app, will serve as a way of opposing the erasure of personal and cultural histories.

I provide several relevant examples of other Locative Media artworks in Section

2.8.11.

2.8 Relevant Works

The original creative work documented in this dissertation is parallel with and indebted to a broad range of artworks in which computer software is granted some degree of agency over music structure. In this section of the Literature Review, I will describe some of the artworks that have influenced my own creative activity, which I frame as a search for a satisfactory transformation of a composer’s agency (beginning with my own) through engaging in code-based composition (as opposed to traditional notational scores). This transformation uncovers new musical languages, performance sites, and ways in which humans may encounter and engage with one another. It is an interplay of agencies, webs shifting in a dance with algorithmic processes automating aspects of music composition and performance. The works described in this section all derive, in various ways, from such automation. I present the works here in ascending chronological sequence.

2.8.1 “Voyager:” George Lewis

George Lewis¹⁵ is an American composer and trombonist who has worked with personal computers since first discovering the technologies in 1977 (Steinbeck 2018, 263). His first computer music work, “The KIM and I,” (1979) facilitated a duet between himself (performing on trombone) and a Moog synthesizer controlled by a program of his own devising (263). The work brought attention to Lewis, beyond the jazz scene in which he was already well-known, and he continued to develop computer applications that interacted with human instrumentalists. In “Rainbow Family,” for example, three Yamaha DX-7 synthesizers were controlled by Apple II computers running multiple algorithms coded by Lewis, that created music in real time while responding sonically to four improvising soloists (264). Lewis calls these computer systems “creative machines,” and sees their capacity to improvise

15. Lewis is a key scholar and artist in the the field of critical studies in improvisation, and has influenced my understanding of agency through his writing and creative work with interactive digital music systems. He is a member of the International Institute for Critical Studies in Improvisation, which hosted a Summer Institute that I attended in 2015. Some of Lewis’ contributions are summarized here: <http://improvisationinstitute.ca/team-member/george-e-lewis/>

as a means to “explore how meaning is exchanged through sound” (George E Lewis 2018, 128). Lewis sees the improvisational capabilities of creative machines as an opportunity to learn: “To improvise is to encounter alternative points of view and to learn from the other; improvising with computers allows us a way to look inside these and other fundamental processes of interaction” (128). Thus, for Lewis, the musicking interactions between human and machine are not empty, but actually facilitate an exploration of ourselves, and our environment.

Perhaps Lewis’ best-known creative machine is “Voyager,” which Lewis first conceived and coded from 1986 to 1988 (George E. Lewis 2000, 34). Like his earlier systems, “Voyager” was designed as an improvising partner alongside which other musicians would perform. He does not see the interchange between “Voyager” and human musicians as a “stimulus/response setup” (34). Rather, humans and machines engage in “a nonhierarchical, improvisational, subject-subject model of discourse” where the software is not seen as an instrument to be played, but as a player itself (34). This quality is reflected in the fact that “Voyager” was quite capable of improvising solo, without needing impetus from human musicians (Steinbeck 2018, 266). Indeed, outside of its capacity to listen, the software only needed two commands from its human user: “start playing” and “stop playing,” which reflected Lewis’ intent that “Voyager” make independent musical decisions, and not act as a mere extension of the human performers which whom it played (265).

This significance of “Voyager,” beyond the technical achievement, is partly founded in what it can reveal about human nature when it encounters computers in shared improvisations. Lewis states that “to improvise is to encounter alternative points of view and to learn from the other; improvising with computers allows us a way to look inside these and other fundamental processes of interaction” (George E Lewis 2018, 128). By imbuing his interactive systems with autonomy and eschewing repeatability, Lewis creates an entity with which we can interact as fellow improvisers. When the improvisation with machines reaches a point where the identities of human and machine dissolve into one, the musical experience is no longer a simulation—it becomes music making, where “a form of artificial life... produces non-artificial liveness” (128). Seen through this lens, “Voyager” and other improvising machines blur the boundary between human and machine agency, and so affirm the value of using interactions with machines as a means to understand ourselves.

2.8.2 “Ambient 1, Music for Airports:” Brian Eno

Brian Eno’s 1978 album “Music For Airports” illustrates the disconnect between the production of algorithmic music and its consumption in linear media such as audio recordings or video (Levtov 2018, 629). The album developed Eno’s concept of Ambient Music, originally introduced in his first solo album, “Discreet Music” (1975)¹⁶. In his Ambient Music, Eno sees music as a support to the atmosphere of a space (Eno 1979). Like Muzak, the specially-arranged popular music arrangements piped into public spaces of the era, Ambient Music is designed to modulate emotions. Eno, however, distinguishes Ambient Music from Muzak, observing the homogeneity of Muzak, how it removes the “doubt and uncertainty” of the source songs, regularizing the environments which it accompanies. By contrast, Ambient Music was a background music designed to induce a calm that encouraged thought, that could sustain concentrated listening as effectively as it could be ignored (Eno 1979).

The recording of “Music For Airports” belies the process by which it was recorded. The linear unfolding of the album is simply a “static snapshot” of the unfolding of one iteration of the piece (Levtov 2018, 629). The music was not scored or improvised by musicians; rather, it was assembled from unpredictable processes, never fully defined by Eno, which could theoretically produce an infinite number of audible variations (269). Indeed, ambient music dismisses distinctions between score, performance, and recording—the first two do not exist in ambient music, and so musicological ideas such as interpretation and authorial intent often associated with Western Tonal music may not apply (Sun 2007, 136). An ambient record’s function is to exist unobtrusively in the background and feel like it has no source, no human ensemble behind it (136).

It is worthy of mention that “Music for Airports” inspired an app called “Air” (Levtov 2018, 632). The app is essentially a digital recreation of the original album (Levtov 2018), where the listener loads the app and hears an endless rendition of the same generative composition. One of the designers of the app, Peter Chilvers, has collaborated with Eno to produce other apps, such as “Bloom,” described in Section

16. Interestingly, the wash of synthesized polyphony in “Discreet Music” was governed partly by a digital recall system (Loydell and Marshall 2017, 2), a relatively early instance of a composer electing to shift their composerly agency to a digital control system. Prior to “Discreet Music,” Eno had experimented with looping using analogue tape, including in his “No Pussyfooting” album, a collaboration with Robert Fripp (of the progressive rock band King Crimson) (Tamm 1990, 58). Both of these looping, textural works seem to embody Eno’s experimental approach to composition, where the composer sets up conditions and then steps back in anticipation of being surprised.

2.8.7.

While none of my works are actually instances of Ambient Music, I have drawn upon elements of the form in my own creative work¹⁷. While “Shards of Memory” foregrounds the music, it is also designed to enhance a physical space and provoke thought. My earlier work “Moonloops” (presented in Chapter 6) has an even closer relationship to Ambient Music. It could generate an unfolding musical wash of sound that was calming and could proceed independently of user interaction, and my choice of instrumentation and texture was informed by some of the same ideals of Eno—to provoke thought unobtrusively while modulating listener emotions. As a snapshot of a larger, unfolding work, “Music For Airports” is a progenitor of dynamic computer music, and, in my own experience, its lilting themes helped me to see the potential in a music that transcends the temporal bounds of the notated score.

2.8.3 “Solo” and “Rhythms:” Joel Chadabe

Chadabe produced two works as original exemplars of his concept of interactive composing, titled “Solo” (1978) and “Rhythms” (1980) (Chadabe 1984, 22). Since the former utilizes an approach similar to my own early computer music experiments, I will describe it in more detail.

“Solo” focussed on the automatic composition of melodies, in response to the performer’s gestures as detected by two proximity-sensitive antennae (22). The right antenna governed the density of sound events, governed by the degree of proximity of the hand to the antenna. Hand motions near the left antenna governed the selection of instrumental timbres through which the synthesizer would play¹⁸. Since agency over musical output was shared between performer and computer, results would typically be unpredictable for both parties. For example, a certain gesture by the performer might trigger the entry of two flutes and clarinets, but the chord that these sounds

17. Indeed, in the liner notes to “Discreet Music” Eno writes “Since I have always preferred making plans to executing them, I have gravitated towards situations and systems that, once set into operation, could create music with little or no intervention on my part.” I believe that my doctoral studies have motivated my slow evolution from a creator of documents to a developer of unpredictable systems mingling agencies, and so Eno’s work has become a personal touchstone, although in my work, I wish to form a dialogue between myself and other agents (e.g. users or listeners) where Eno in these early works engages with the machine in a way that does not involve the listener in the construction of the musical work.

18. For a recording of “Solo” and further information, visit Chadabe’s web site at <https://joelchadabe.net/solo/>

form would be unpredictable, as the computer has agency over harmony since it is in control of pitch selection (Chadabe 1984, 25). Through the dialogue that this interaction supports emerges a unique musical structure over which both computer and performer might claim authorship. Chadabe himself saw in the early development of interactive composing a redefinition of the roles of composer and performer (26). In particular he notes that the tasks involved in interactive composing relate to defining an algorithm and choosing which variables should respond to the control of the performer and which remain in the domain of the computer's agency. In a traditional score, the composer understands the grammar of music and its causal effects upon instrumental output when interpreted by an instrumentalist. In interactive composition, however, outcomes are more nebulous, and musical structures that emerge are "the result of the specific functioning of the system in a particular performance," with the full composition being revealed only after all parts have been generated through the joined action of performer and algorithm (26).

2.8.4 "Sonic Meditations:" Pauline Oliveros

To Pauline Oliveros, listening grants the listener the possibility of a psychological transformation (Jacquin and Polverel 2020, 1). She describes her concept of active listening, which she terms "Deep Listening" as "listening in every possible way to everything possible to hear no matter what you are doing.... Deep listening represents a heightened state of awareness and connects all that there is. As a composer I make my music through Deep Listening" (Oliveros 1999, 1). As her compositional practice evolved, she abandoned the new music paradigm that enforced rules and an entrenched the distinction between performer and audience, returning to what she termed "Sonic Explorations," inclusive forms of musical sharing (Oliveros 1999). Here I will briefly introduce her series of works entitled "Sonic Meditations," collected in a publication by Oliveros in 1974 (Jacquin and Polverel 2020, 2).

"Sonic Meditations" originated in Oliveros' work in improvisational performance, for which she formed the "Women's Ensemble" at the University of California in the early 1970s (2). The ensemble featured both vocalists and instrumentalists, and initially focussed on musical forms incorporating the slowly-developing timbral variations associated with Oliveros' early musical approaches. Over time, however,

the members of the ensemble began to allow involuntary changes to feature more prominently in their performances—that is, the music became more a product of unconscious efforts than conscious ones (Jacquin and Polverel 2020, 2). This performance paradigm required a transition from production to listening, as the unfolding of unconsciously-produced music required attentive ears. Oliveros describes “Sonic Meditations” as an attempt to return the control of sound to the individual alone, and within groups especially for humanitarian purpose” (Oliveros 2017, 2). The scores are textual, designed not as guides in the formation of a musical performance, but as “attentional strategies” to help the participant listen to their surroundings (Jacquin and Polverel 2020, 2). A performance of one of the “Sonic Meditations” can be joined by anyone regardless of musical skill or training. Participants in a performance engage in a collective form of Oliveros’ Deep Listening (3). Oliveros held that consistent, sustained engagement in such activities is a form of meditation that offers healing benefits (Oliveros 1999, 3). Indeed, Oliveros felt that listening is a foundational human capability. She states “How we listen creates our life. Listening is the basis of all culture” (3). So, the textual scores of the “Sonic Meditations” ask questions and challenge the participant to think about their present and past sonic environment—for example, asking the participant to come up with a sound that reminds her of home, and share it with a neighbouring participant (Oliveros 2017, 2).

Oliveros’ work is relevant to my work in that it shows a path towards effectively decentralizing agency. Oliveros is not interested in imposing order upon surrounding acoustic space. Rather she envisions an inclusive form of music making that emerges from communal action, memory, and will. I find her vision of music intriguing partly because, like myself, she has relocated the site of her capacity for musical production from notation to another domain—she uses textual instructions as a means of inciting communal sound production and awareness, and I cannot help but recognize that her “Sonic Meditations,” in their flow of instructions, resemble my own coded scripts. Consider the following excerpt from “Sonic Meditations”:

Meditation XXII

“Think of some familiar sound. Listen to it mentally. Try to find a metaphor for this sound.

“What are the real and imaginary possible contexts for this sound? How many ways does or could this sound affect you? Or how do you feel

about it? What is its effect upon you? How can this sound be described?

“As a group meditation; sit in a circle. Find a sound common to all, then ask the above questions one by one. Allow plenty of time between each question. When all of the questions have been asked, the group shares their answers.

“Variations: Try the same meditation with

1. an imaginary sound
2. a live sound
3. a remembered sound.”

(Pauline Oliveros, quoted in Jacquin and Polverel 2020, 6)

“Meditation XXII” is an algorithm, guiding the agencies of the participants, but not constraining them in the manner of a typical Western tonal notated score. It is a process that she has specified, and in which participants choose willingly to engage. The goal is not the performance, but rather is partly framed as the “enhancement and development of aural sensation” (Oliveros 1984, 141). In my own interactive music I also seek to promote awareness, particularly of empathetic connections between participants. The common use of algorithms to promote awareness is a method that I share with Oliveros, and it points to a nascent potential in algorithmic art to shape deeply its participants’ experience.

2.8.5 “Experiments in Musical Intelligence:” David Cope

David Cope, first known as a composer, constructed a series of generative music systems known collectively as his “Experiments in Musical Intelligence,” or “EMI” (Cope 1992). Cope’s programmatic creations are often tasked with the emulation of specific composers’ styles, including those of Cope himself and of famous classical composers such as Bach and Mozart (Cope 1999, 79; 1992, 69). EMI was capable of analyzing databases of any individual composer’s scores, not just those of Cope himself; it can identify recurring features or “signatures” in a given composer’s music that, Cope argues, have “stylistic importance.” EMI then uses these structures as guides to create “new instances of music in that style” (Cope 1999, 79).

The original music produced by EMI has, at times, confounded critics by prompting them to question whether music is an exclusively human construction. For example, Douglas Hofstadter, a cognitive scientist whose research includes study of consciousness and the nature of artistic creation, played through an emulation of a Mazurka (arranged for piano) in Chopin’s style, and, with “mounting confusion and surprise,” began to wonder that the piece “seemed to express something” (Hofstadter 2001, 4). He was “shaken” by this occurrence, as this revelation posed the uncomfortable suggestion that computers were removing the “profundity of the human mind’s sublimity... [which] seemed somehow humiliating, even nightmarish” (6).

I share these concerns, though I have come to embrace algorithmic music as a doctoral student. My own understanding of my role as a composer or artist actually resembles Cope’s, since he comfortably inhabits roles of both composer and programmer. Yet, Cope might be seen to threaten the humanity of music itself with his Frankenstein-like generative creations, leading us to call into question our assumptions regarding the necessity of human creativity, or human expression, within music. Yet, before positioning computer-based algorithmic composition as an unprecedented occurrence, it should be noted that algorithms are not unknown in the musics of earlier eras. Many familiar formal constructions specify a template for a work—the fugue a staggered pattern of interleaved melodies, the sonata allegro form a sequence of sections with specifications governing the configuration of their content. These conventions impose a template-like process on composers, and provide a framework for listeners that may partly govern their expectations. If an algorithm is a set of rules that, when implemented, solve a problem, then rules of counterpoint, voice-leading, or harmonic modulation have an algorithmic character. David Cope has gone so far to argue that “all composers are algorithmic composers” (Cope 2015, 405), proposing that composers and listeners alike rely on a (possibly unconscious) common understanding of algorithms governing the progression of music that allows them to predict its unfolding (410). Once one hears Revel’s “Bolero,” for example, his ruleset may seem obvious—structure the work as a crescendo, introduce additional instrumentation to continually magnify its mass, alternating two themes throughout. Understanding that the piece is a crescendo (whether one is familiar with the term or not) can be a source of its effect, for it grants the piece coherence. In broader terms, if we lacked such algorithms, it may be hard to see how music could communicate

at all—the process of its unfolding would appear ungoverned, wilder than the freest improvisation. Algorithms provide order and structure, and so a common ground upon which composer, performer and listener can meet.

Hofstadter reeled upon hearing EMI's new Chopin-style mazurka, because the music seemed to him to express something. Once he experienced emotion while performing EMI's music, the question as to what was being communicated by EMI was raised. Hofstadter reported that had he been informed the piece was composed by a human, he would not have doubted its expressiveness, as it was “not emotionally empty” (Hofstadter 2001, 6). This meaning is not an unprecedented occurrence for music-generating algorithms, for the reasons stated above. What is significant is the form of connection between those algorithms and the composers involved: Chopin was present as an agent in the form of his scores, and Cope as the coder. I feel that by passing this sort of emotional Turing test, EMI does not detract from music, but extends it. Cope has constructed an algorithm that, in the interactions between Cope the coder, the composer who provides the library of scores, and the performer who performs the score for listeners, makes music that raises unexpected emotions in Hofstadter.

This is as it should be, for musicking, whether bringing together people mediated by the codes of notation or a programming language, is about the people who participate. When Hofstadter felt emotions, he was disconcerted, arguably because he felt empathy—the sense that a music student had written the piece as an exercise. There was no student, but rather an algorithm. It was like embracing a machine, and then realizing that you are actually alone. But musicking is not a lonely act, for its strength is in the building of relationships. If Hofstadter felt empathy, I believe it was genuine. After all, if Foucault is to be believed, we all are separated from reality by the filter of language. The emotions Hofstadter felt were his own construction, and what he perceived was as real enough to him to be meaningful. He was not connecting with the Author/Composer, but the impetus for the event was still human, issued from a network of agencies—those of Cope, Chopin, and he himself—that created a potential for feeling. And he felt something. Perhaps that was enough.

2.8.6 “Aspects of Gaia:” Roy Ascott

As described in Section 2.5.4, new media artist Roy Ascott is known for casting art-making in cybernetic and telematic terms, where artist-controlled determinism gives way to feedback loops and a creative interplay between multiple agents. The physical object—a painting, a bound text, or a notated score, perhaps—has weight and substance. The computer, even during the time when Ascott’s earliest articles were published, has been challenging the solidity and inertia of such physical art objects. Ascott’s own creative work has moved between these forms. An example is his 1989 work “Aspects of Gaia.” He draws for inspiration upon James Lovelock’s “Gaia Hypothesis,” a theory in which the Earth is understood as an organism unto itself, where its living components (the organisms that live there) and non-living features (such as climate, atmosphere, and geography) have formed as they are in a way that maintains the Earth’s vitality¹⁹. The work as an example of a “Gesamtkunstwerk,” which may be defined as a “synthesis of the arts, human, and machine into a whole” (Rafie 2016, 1). The cybernetic quality of the work manifests in the form of audience participation in the construction of the work. There were two sections of the installation, which was installed in Brucknerhaus, Linz, Austria. One in an outside area featuring horizontal monitors displaying streaming texts and visuals on the theme of Earth’s ecology. Additionally, there was a trolley running under the site that patrons could ride, allowing them to see LED lights spelling messages on themes related to the ecological content of the upper installation.

The work was participatory in that viewers could modify content as it was transmitted across a network to remote locations, and this content could be collected by the computers running the work and used to form new content (2). In this way, audiences could participate in the construction of the work, and the boundaries between spectators, designers, content creators, and the art object itself became porous (2).

Of the things that interest me about “Aspects of Gaia,” its conception as a second-order cybernetic system is particularly relevant to my work. Cybernetic systems feature a feedback loop between two entities, potentially machine and human. In first-order cybernetic systems, the focus is on “autonomy and regulation,” where the

19. This description of the Gaia Hypothesis, and documentation of Ascott’s associated work, is partly based on an article at <http://https://artelectronicmedia.com/artwork/aspects-of-gaia-trolley-under-brucknerhaus-with-networked-led-messages/>

system functions according to a discrete function (Herber 2014, 373), such as the maintenance of a consistent temperature in a thermostat. A simple thermostat generally has two components, a thermometer and a heat source, along with a single goal, namely the maintenance of a consistent temperature. When the ambient temperature falls below the threshold, the thermometer detects it and activates the heat source. When the temperature rises above the threshold, the thermometer triggers the heat source to stop. In music, a first-order system might be as simple as a volume control knob that the user continually modulates to keep loudness at a comfortable level, or, in a video game score, a process designed to increase the tempo of the score when the player’s avatar is running. In a second-order system, there are multiple variables, allowing for multi-dimensional changes. In such a system, the human “observer” of the system is confronted with a multiplicity of effects for a single interaction, and can exert a different quality of control over the system (374)²⁰ . An example might be a generative music app such as “Bloom” (see Section 2.8.7 for a discussion of this app), where taps on the smartphone screen impact multiple musical dimensions of subsequently-generated tones, including timing, pitch, and loudness.

In “Aspects of Gaia,” Ascott builds on the observation that the Gaia Hypothesis actually portrays the Earth as a self-regulating cybernetic system essentially functioning as a feedback system to maintain its condition as a planetary homeostasis²¹ . “Aspects of Gaia” is a second-order system because the observer who encounter the work are not constrained by the function of the system. A second-order cybernetic system is a *social* system, according to Heinz von Foerster, the originator of the concept of the second order cybernetic system. Foerster states of such a system that an “observer who enters the system [is] allowed to stipulate his own purpose: he is

20. Arguably, such feedback loops may also be observed in digital music-distribution apps such as Spotify or iTunes, where algorithms suggest new tracks to listeners based upon prior listening habits, and these new tracks, once added to users’ listening queues, eventually become fodder for the next set of recommendations. Such a situation raises questions of how the agency of users of such music distribution systems might be curbed or magnified by their interactions with such a cybernetic system, and how, through such interactions, their musical preferences might (perhaps unknowingly) be transformed.

21. Homeostasis refers to the regulation of bodily conditions by feedback-based physiological systems. Homeostatic systems ensure that the body can sustain its functioning, for example by maintaining blood oxygen in an optimal range, or keeping the levels of adrenaline at a value proportionate to the organism’s current optimal state of repose or alertness.

autonomous” (Von Foerster 2003, 286). This quality distinguishes first- and second-order systems. In the case of the former, the system’s function is pre-defined, and admits no other purpose: a volume knob will only affect loudness and a thermostat only offers control of ambient temperature. In a second-order system, the observer enters the system to use it according to their own goals, and the system supports such intent. Ascott’s “Aspects of Gaia” was innovative partly because it was a second-order cybernetic system: a participatory system sustaining social interplay in which those who encountered the system could act according to their intent (Rafie 2016, 2).

In my perspective on agency in algorithmic music, a second-order cybernetic system is essential. Ascott’s “Aspects of Gaia” is one of the first examples of a digital work that created a social space in which participants could retain and project their innate intent through the work. The arc of my doctoral projects displays a model of interactivity that begins with no interaction (my notated scores such as “We, the Artificial”) through first order systems (such as “Q-Chords,” which has an extremely specific function from which a user is not able to substantially diverge), to my final work, “Shards of Memory,” which I argue approaches a second-order system, in the sense that it grants the user some influence over how the system responds to emotions projected by elements of the work (more detail is provided in Chapter 7). Ascott’s own work embraces a multitude of agents and grants them considerable freedom to act self-expressively. That the work is not specifically centred on music does not take away from its relevance here, because it is an instance of a social convergence that fosters interpersonal interplay. As such, its objectives align with my own, and represent an early application of digital technology towards the goal of magnifying empathy and respect that demonstrates how fully second-order cybernetics can be model for promoting the social harmony that I aspire to facilitate in my own creative work.

2.8.7 “Bloom:” Brian Eno and Peter Chilvers

Brian Eno and Peter Chilver’s iOS app “Bloom” (2008) was the first interactive music app I encountered as a doctoral student. It is identified by Guedes as an instance of RTC designed for “common/lay users” (Guedes 2017a, 2). Chilvers described the app as having being derived from Eno’s work on the video game “Spore”

(Chilvers 2016), a game which relies heavily on generative processes throughout its design (K. Collins 2009, 11). Brian Eno’s score was prototyped in Max/MSP and then realized for the published game in Pure Data (11), both of which are coding environments designed for manipulating music data, as audio recordings or as midi-like events, as documented by the coding environments’ designer, Miller Puckette (Puckette 2009).

“Bloom” is notable for the way that the agencies of its designer and its users may encounter one another through the app and its musical output, which its designers describe as “part instrument, part composition, and part artwork” (Lowgren 2009, 7). Using the touchscreen of the iPhone as a control surface, users tap the screen to evoke and alter patterns of musical tones with their pitch and timing influenced by the algorithm’s analysis of the positioning of the taps (7). This interaction makes the user into something like a musician, but not entirely, in that, unlike performance upon a conventional instrument like a violin, the “Bloom” app can act in a partially autonomous manner, generating tones that automatically conform to key, timing, and timbral constraints which maintain the expressive quality of the overall music²². The automated interpretation of users’ interactions with the app means that the agencies of those users contacts the agencies of the app designers, and the music that “Bloom” generates owes its form to both. Bloom is one of those apps that challenges established musical roles by shifting the aspects of musicking governing sound structure between such roles. This blurring has been a key concern of my doctoral research, because it represents a transformation of agent roles in music. As one of the first generative music systems I encountered upon enrolling in my doctoral

22. It is important to note that, while the interaction model for “Bloom” involves an encounter of agencies between user and composer/programmer, it is also possible for generative music system designers to design generative systems for their own use, which might be conceived as a “closed” agentic system. The English IDM duo Autechre (Rob Brown and Sean Booth), for example, used the Max coding environment to create generative systems that they could manipulate live in performance—effectively jamming with the system—as well as in the studio (Tingen 2004). While Autechre’s audiences are not given an entry-point through which to interact with the system, the bespoke patches Autechre designs does externalize some of their own compositional agency into algorithms, which then feeds back to them. While the creative work I document in this thesis involves a sharing of agency with its audience as a means of fostering inclusion and interpersonal contact, Autechre’s approach does suggest the value of algorithmic creativity in situations where the composer seeks novelty and surprise, but chooses to exclude their listeners’ direct input. This approach to coding is the mirror of my own goals with my doctoral research, looking inward instead of outward, and in a sense it provides a glimpse of a path I might have chosen in my development as an artist, if I had chosen to continue to elevate my composerly agency, rather than mingling it with others’.

program, it provided an early impression of a construction that is part artwork and part instrument, a model of bringing the audience for an artwork into the processes of its creation that has grounded my conception of interactivity in my own doctoral studies from the beginning²³ .

2.8.8 “The Cave Of Sounds:” Music Hackspace

Building upon Christopher Small’s idea that collective musical activities help participants to form a common identity, the “Cave of Sounds” is an interactive installation premiered at the Barbican in London in 2013. The project was developed by a team of volunteers from the “Music Hackspace,” a community of individuals exploring “new forms of music, sound art and musical practice made possible by technology.” The project is the outcome of Tim Murray-Brown’s residency with the group (Murray-Browne et al. 2014, 307-308).

The installation consisted of eight digital musical instruments, each designed by a different volunteer (including Murray-Brown). Recognizing that designing a musical instrument can be a process of self-expression, the digital musical instruments featured in the “Cave of Sounds” were designed not only as individual expressions, but were conceived and designed together specifically so that the group could explore the similarity between composing instruments and performing with them in improvisational settings (307). The instruments were designed using many different technologies and forms of interaction. For example, three instruments made use of Kinect motion sensors, others used Max/MSP and its Ableton implementation Max4Live, Arduino programmable circuit boards, traditional instruments such as a flute, and sound sources randomly collected from the internet (308-309). When installed, instruments were arrayed in a circle, and were then available for the audience (patrons and

23. More recent apps have further innovated paradigms for interactive music app. Endel, a generative music app notable in that it was signed to Warner Music, for whom the system is delivering twenty albums, with a proportion of the royalties going to its programmers (Makhmutov, Varouqa, and Brow 2020, 3003). On Endel’s official website (<https://endel.io>, the company touts its use of smartphone sensors to deliver custom mobile musical sound environments to users as they go about their daily activities. Where “Bloom” allows inout through its users’ screen taps, Endel monitors more sensor streams, including location, time of day, weather, and heart rate. I regard Endel’s use of these sensor streams as a significant development in algorithmic music, since such data can help the app to form a representation of its listeners which, in turn, can help the app to produce music that engages those listeners, whether they are conscious or unconscious of the music-generating feedback loop in which they are engaged.

other passers-by at the Barbican) to engage with them. (Murray-Browne et al. 2014, 307)

One contribution of this project that is particularly relevant here is how the creators of the instruments frame the complex web of agencies that are active within the installation as it is performed. There are two broad categories of agencies at work: those of the instrument designers, and those of the audience/performers. Each instrument is the product not just of its creator’s personal practice, but also of the process of participating with the other designers in creating the ensemble as a whole (308). The balance between these different voices is experienced by audience members who perform on the various instruments, while offering these performers a chance to connect with others who share the space playing different instruments (308). Thus the “Cave of Sounds,” like many efforts within the NIME community, points to a music where boundaries between composer, performer, instrument designer, programmer, and audience are blurred (310), disrupting established hierarchies of musical role enforced by *Werktreue* and the institution of the classical concert, in favour of a more egalitarian understanding which values the formation and strengthening of interpersonal bonds.

2.8.9 “Musebots:” Arne Eigenfeldt

Arne Eigenfeldt understands real time composition in slightly different terms from Guedes’ description discussed above. Eigenfeldt defines real-time composition as “the application of musical agents to interact in musical ways, during performance” (Eigenfeldt 2011, 146). In his applications, intelligent agents share control over output music. Each agent is autonomous, has inbuilt knowledge of how to interact to with its fellow agents, as well as the capacity to control a musical gesture, which might be formed of elemental musical properties like pitch or timbre (147). Music emerges from the social interactions between the agents, and the user’s control of the music is not accomplished by altering low-level musical parameters, but rather high-level parameters governing the agents’ behaviour (148). Some interactions, Eigenfeldt says, can lead to musical events that seem “magical,” arising serendipitously from improvisation shared between agents (148). His work with these musical environments populated with interacting autonomous agents continues with the Musebots.

Eigenfeldt defines Musebots as “pieces of software that autonomously create music collaboratively with other musebots” (Eigenfeldt et al. 2017, 1). They are consistent with the broad definitions of a bot provided earlier in this chapter—they are autonomous, cooperative algorithms with specialized functions²⁴. The ensemble can be assembled from available Musebot types, including beat generators, harmony generators, and several other types, all of which communicate through a special “Conductor” bot, which provides synchronization through saving as a hub for the other bots’ messages (2).

An interesting recent example of a work involving Musebots performing in collaboration with human performers is “Unauthorized,”²⁵ a collaboration between dance/theatre artist Kathryn Ricketts, Eigenfeldt, and the Musebots. Presented in October, 2019 at Simon Fraser University in British Columbia. The work features Ricketts performing as the clown Rufus, who “struggles to find humour in dissonance”²⁶. The introductory movement is accompanied with music from Eduard Greig’s score to the play “Peer Gynt,” during which time the Musebots analyze Ricketts’ motion in real time to determine their approach to orchestrating the remainder of the performance. Once they begin to sound, the Musebots sense Ricketts’ movements and use this data to govern parameters of their musical performance, to which Ricketts can respond, create a feedback loop between her and the Musebots. The individual sections of the work each feature a different staging and feedback dynamic, forming distinct explorations of interaction.

Musebots are social entities, as the music they produce is shaped by multiple Musebots acting collaboratively. In “Unauthorized,” Ricketts herself becomes an agent, performing alongside the integrated agent formed by the Musebots acting in a totality. The performance emerges from a hierarchy of agencies, Ricketts encountering the collectivized Musebots, an assemblage of the individual Musebots’ agencies, through which Eigenfeldt projects his own agency in automated form. Of all of the artworks in this section of the Literature Review, “Unauthorized” features perhaps the richest encounter of agencies, between performer, programmer, composer, user,

24. A demo of the Musebots can be seen at <https://www.youtube.com/watch?v=Rb2n9L-Jtjo>

25. A video of the performance can be viewed at <https://www.youtube.com/watch?v=XTJbMWGM13w>

26. Eigenfeldt’s description of the work can be viewed at <https://aeigenfeldt.wordpress.com/unauthorised/>

and autonomous algorithmic entities, a quality I find inspiring.

2.8.10 “Threnoscope:” Thor Magnusson

The Threnoscope is a musical work that involves live coding, which conceives program code as a musical instrument. I find the form particularly interesting as it is a practise where the roles of composer, programmer, and performer are heavily blurred. I will give a brief description of live coding, then introduce the work.

Live coding is an artistic form involving programming on stage with interpreted languages as a performance²⁷ (Blackwell and Collins 2005, 3). Live coding is a practice merging “real-time composition and performance in a participatory context” according to Thor Magnusson (Magnusson 2014, 268) and uses the form to find points of contact between composition and coding. In this comparison, one may find evidence supporting the unprecedented nature of computer music, arguments that suggest that the artist’s computer has capacities that transcend the scope of many other technological innovations.

He draws parallels between coding and traditional Western art music score writing, noting that “computer code is a form of notational language” just like musical notation. (268) Notation specifies events through a systematized set of instructions that form a code that can be read by musicians and computers alike to render its performance. (269) For composers, notation serves as a mnemonic device that compartmentalizes composing processes and aids the development of complex scores. (269) As written, WAM notation typically specifies events progressing in a linear manner on a timeline, a sequence of notes and other symbols that define the linear unfolding of the score. Code, on the other hand, is different—Magnusson calls code a “peculiar form of notation and computers... peculiar interpreters.” (270) First and foremost, computers can appear to transcend the linear imperative of the clock, by moving, delaying, or surprisingly initiating events on a timeline that can defy the rigidity of the notated, linear score. (270) While code itself may follow a linear path, its “playhead” (Magnusson’s tape-based metaphor for the ever-shifting point of code execution) can jump through the code in a nonlinear fashion that may be dictated in part by the code itself, influenced by whatever inputs that code is designed to admit. In observing

27. Interpreted languages are programming languages where procedures can be executed directly without need of compiling

this capacity, the beginnings of a divergence between score and code start to appear.

Magnusson introduces the Threnoscope²⁸, a music-making environment that features a circular visual component that maps onto the sound output, consisting of code-governed drones. (Magnusson 2013, 2) When Magnusson launched the project, he was interested in how it would be perceived, whether as an environment for live coding, or as being essentially a singular musical work. The system poses certain constraints, such as its focus on drone sounds, its visual interface which guides its use, and the set of commands available to its users. Magnusson wonders if its constraints (as solidified in its coded design) would appear to mirror those of a notated score, which, in performance by a musician, may specify events but not necessarily their entire mode of expression. Elements of a musical performance not specified in a score are widely regarded to be essential to its capacity to satisfy, given the long-standing observation that utterly literal performance renderings of a composer's score lack expression normally in the purview of the interpreting musician (Repp 1989, 243). In live coding, there is a tension between those who have a strict deflection of live coding, which requires writing and changing of algorithms in performance, and those who accept a more inclusive view where, for example, reassigning variable values in a prewritten program is enough. (Magnusson 2013, 1) To me, this tension mirrors my own questioning of my agency in music, whether as a composer who relies on externalized, computer-based extensions of my mind, or as a composer using program code to automate his style of compositional action, giving up control over musical structure in an attempt to innovate its form.

2.8.11 Works of Locative Media

The following examples represent works of locative media relevant to my studies.

“murmur”

“murmur” (begun in 2002) was a locative art installation project that curates stories in specific urban locations (Wershler 2008, 407). Conceived by Gabe Sawhney and Shawn Micallef during a residency at the Canadian Film Centre's New Media

28. A demo of the Threnoscope can be found <https://vimeo.com/63335988>

Lab (Wershler 2008, 407). The project combined elements of locative media and site-specific performance, two very different disciplines that share in common a focus on place and context as sources of meaning (Eaket 2008, 31). The installation incorporated cell phones as a delivery medium and distinctive green signs, labelled with a numeric code, as a marker. Travellers through the city would see the sign, and then either call “murmur’s” phone number on their cell phone, or tap on the sign as found in “murmur’s” web interface (Wershler 2008, 407). After entering the code specific to the sign, the traveller would then be played an audio recording by a “storyteller,” a (possibly anonymous) contributor of a story that relates to the place where the traveler is currently situated (Eaket 2008, 33).

The installation uses the intimacy of storytelling and the grounding of space to collapse the social distance between strangers (34). At the same time, it introduces an alternate conception of the space the traveller finds themselves in, that of the storyteller. The traveller must “reconcile” these two views of the place, and in so doing form its meaning—that the place is “both material and historical, as well as deeply social” (38).

While the work is speech-based, rather than musical, the goals of the installation are aligned with my own understanding that sharing stories can be the beginning of an empathetic connection. The traveller seeks the signs to find the storyteller. The storyteller reveals her experience. The traveller reconciles the view of the storyteller with his own understanding of the place. The last act is an act of empathy, of seeing from the storyteller’s perspective. If the storyteller is honest, and has something important to say, I believe that an empathetic connection can be formed. There is little distance in this form of sharing from Small’s musicking, which is powered by sharing the experience of musicking, and is an aim of my own creative work documented in the following chapters of this dissertation.

“Pokemon Go”

“Pokemon Go” was introduced in 2016 as a smartphone gaming app. The game belongs to the Pokemon universe, which envisions a world where humans live alongside wild creatures (the Pokemon), whom the humans collect, train, and battle (Bainbridge 2014, 399) The Pokemon were conceived in 1996 by Satoshi Tajiri, a Japanese video game designer, and through the frequent introduction of new video games,

films, and television shows, the Pokemon Universe has expanded to become one of the most successful game-based franchises in the world, and a popular example of trans-media storytelling (Bainbridge 2014, 399). The “Pokemon Go” game extended this universe into the realm of location-based augmented reality, by superimposing upon the real world animated Pokemon that could be captured by players through a gesture on the smartphone screen simulating the throwing of an in-game device that would “catch” the Pokemon and add it to the player’s inventory. Through displaying the forward-facing smartphone camera feed on the screen and superimposing the animated Pokemon upon it, the game turned the phone into a window through which Pokemon could be discovered and caught. The Pokemon and other in-game objects were instantiated in fixed, real-world GPS coordinates (Paavilainen et al., n.d., 2494). The game enjoyed immense popularity upon its release, and was recognized as the “first meaningful experience with location-based gaming for tens of millions of people around the world” (Colley et al. 2017, 1179).

While the designers of “Pokemon Go” did not incorporate complex, GPS-aware algorithmic music generation systems, the game has been seen as fostering social interactions through a game design that encourages players to engage with one another. The augmentations of real space can support interactions between multiple players, friends or strangers, in the same space, through, for example, the formation of local teams of trainers, or sharing the benefits of luring rare Pokemon to a space (Paavilainen et al., n.d., 2496). I see in this formation and reinforcement of social bonds a certain compatibility with musicking, and wonder what music the Pokemon might make if they were given the opportunity.

“The Rough Mile, Part 1”

“The Rough Mile” (2017) is a site-specific, locative walking experience²⁹ designed for pairs of friends (called “visitors”) to complete jointly (Hazzard et al. 2017, 1). Each visitor in the pairing begins and ends their encounter with the app together, but the walk itself is taken solitarily. “The Rough Mile” incorporates a spoken narrative accompanied by music and ambient environmental sound. The narrative is scripted as a first person account of a challenging time. Two human actors play characters at

²⁹ A video documenting the project can be found here: <https://www.youtube.com/watch?v=RtRmFIPVCIY>

specific locales along the path, approaching and interacting with each visitor while recoding the audio of the interaction. The recordings contain each user's choice of songs that they believe their friend would like to hear at particular points in the narrative, as well as their reflections on their choices (Hazzard et al. 2017, 2). In "The Rough Mile, Part 2," these audio recordings and song choices are used as musical material, but my main concern here is the sound in Part 1.

The music soundtrack of "The Rough Mile" serves in a subordinate role similar to the score of video games and films (2-4). The visitors walk a fixed trajectory through the space, and audio events are conceptualized as "scenes" which are triggered along the path by the visitor moving to new locations (4). Audio layers were associated with narration (one-shot audio recordings triggered by waypoints or time), ambient sound (presented through bone-conducting headphones), and music which could be mixed as the app functioned from multiple stems. The music was intended to contribute to the affective experience of the narrative, colouring the visitors' perception of the environment (4). At the same time, the music was able to provide an informative function, such as marking the global boundary that encapsulates the site.

Some visitors who experienced the app reported that it changed their understanding of the site (6). In some cases, the app drew their attention to features of the environment they had not previously noticed, whereas some also noted a sense of detachment in the experience. The designers considered their soundtrack successful in part because it was seamless and not distracting (6). Rather, the layers "combined to create one thematic contour," where music "services the spoken narrative and the physical location in the same manner as those music soundtracks found in films or computer games" (2).

To me, this work forms an interesting mix of agencies, introducing live performers into the interactive app environment, while using sound as a guide to the visitor, prompting them along their narrative. Yet the music in this work is subordinate to the other elements, rather than taking its own space. That decision is perfectly reasonable in the confines of the app aims; but from my perspective, the choice of restricting the music to a state of "minimal activity, constituting a sustained drone" (4) would become a missed opportunity to introduce more of a composer's agency into the experience.

“Proprius”

“Proprius” (Ozcan and Camci 2018) is a 2018 augmented reality composition that uses the sonification of organism behaviour as it occurs in a natural ecology (1). Sonification can be considered a form of computer-based display that is auditory in nature and conveys meaning, such as the click rate of a Geiger counter indicating the relative amount of radiation and the phone ring indicating that someone wants to talk by phone (Hermann and Hunt 2004, 5). The score of “Proprius”³⁰ uncovers the behaviour of organisms through musical sonification, for example by indicating the behaviour of birds by representing them with an FM synthesizer, modulating carrier frequency in proportion to birds’ size and modulator frequency to the birds’ energy (Ozcan and Camci 2019, 2-3). The listener is brought into the composition as a “disease agent,” who, when in proximity to an organism, causes their health to degrade and compels them to try to move away as an escape (5). Compositional decisions were inspired by ecological relationships and draw attention to interactions between the listener and artificial agent. The organisms form a musical ensemble whose performance reveals principles of ecology, in which the listener becomes an agent in its creation.

I found the abstract nature of the score intriguing, once I learned to identify how its form maps onto the behaviours of the ecosystem. I couldn’t help but wonder if the use of a tonal music system, such as might be found in a Hollywood film, might be able to help listeners understand the scenario of the app, using their innate understanding of tonal music conventions to support the sonification. For example, birds are commonly identified in music, such as the chirping birds in Mozart’s opera “The Magic Flute” or the many works of Messiaen that draw on bird song, such as his “Reveil des oiseaux,” of which he said “There is nothing but bird-songs in this work” (Hold 1971, 116). We may think of a work in a physical space as a “soundscape,” Canadian composer R. Murray Schafer’s identification of the auditory analog to “landscape,” discussed in Section 2.7.10. Perhaps that model of composition seems suited to AR music, given the importance of place. My own view, however, is that AR spaces may be scored abstractly or not, and the sonification of elements of the space (real or virtual) can use tonal music effectively, leveraging the listener’s enculturated understanding of music

30. Audio recordings of the work can be heard here: <https://soundcloud.com/zeynepozcan/sets/proprius-audio-examples>

and its relationship with space, an approach I explore in Chapter 7.

“Pharos AR”

Released as an app for Android and iPhone in 2019, “Pharos AR” is a continuation of musician Childish Gambino’s “Pharos” series of events (Childish Gambino is the stage name of actor Donald Glover), produced by MediaMonks. In previous events, Gambino had experimented with extending audiovisual norms of music festivals with immersive, technology—driven experiences, such as the 2018 Pharos festival which used projections of interactive visuals and audio conceived in the Unreal video game engine (Miller and Crawford 2019, 2). Gambino followed this spectacle with another one, this time in the form of augmented reality constructed in Unity. The artwork is also a commercial product, and is not yet well documented in the scholarly literature. This description is informed by my own use of their app, and media and promotional material from Unity, MediaMonks, and Gambino himself.

Pharos AR is an artwork and a multiplayer social experience centred on Gambino’s music. The app user scans their immediate surroundings and the app maps the surfaces, allowing it to procedurally generate a visual augmentation of the user’s surroundings, seen through the phone screen in the manner of a first-person shooter video game. Movements through physical space are tracked, to synchronize it with objects in the game world. Once the cavern is formed, the user finds themselves in a cavern mapped onto the physical space they occupy. The user identifies glyphs suspended in the AR space by pointing the forward-facing camera at them, which are beamed off the wall and animated as virtual dancers. When enough have been located a portal appears that leads the user into an animated sequence of events in which the user is sent upwards towards a black hole, through the sky and past moons. The trip is a visual accompaniment to Gambino’s song “Algorhythm.” Once the black hole is reached, the user returns to her surroundings, only to discover that some of the dancers have made the trip with her, and are now dancing on the floor. At that point, the experience ends, with little additional content to encounter.

Pharos AR can certainly be seen as a somewhat confining experience. The narrative is unchanging between play-throughs, and apart from the virtual mapping of the room and the exploration for glyphs, the events follow a sequence that is quite linear. What I love about Pharos AR is the spectacle of it. The music, imagery,

and narrative all are tightly integrated and form a world of sumptuous beauty and energy. Unlike a typical video game score, the imagery and music seem to be of equal importance, placing the user in a simulated sci-fi electronica concert. I personally feel that, even though the experience is brief and rigidly-sequenced, Pharos AR points to imaginative future AR applications where the experience converges with popular cultural forms like the sci-fi blockbuster film, epic video game, or electronic music festival, while placing the player within its world as an active agent in that world's construction. To accomplish that, I feel Pharos AR would have needed to extend its length and its model of interaction. Nevertheless, the app is a tantalizing glimpse of possible future applications of mobile and AR music.

Chapter 3

Interdisciplinarity And Methodology

Whether the software designer wishes it or not, software unfailingly projects his or her mentality and personality..... When you use software you are spending time with its creator.

(Miller Puckette, in (Puckette 2009))

3.1 Introduction

According to J. Gary Knowles and Sara Promislow, “knowing through the arts is more than mere knowledge about the arts.” (Knowles and Promislow 2008, 10) The methodologies through which this thesis has emerged are energized by this idea, that through making art, whether in musical, visual, or programmatic media, I may enrich my understanding of my personal artistic role, and uncover new ways of being a person who creates and shares music. One criterion by which I judge whether one of my artworks is successful involves recognizing whether those who engage with that object change in some way, in terms of their immediate cognitive state (emotional and intellectual), or in the more persistent nature of their social connections and attitudes. While I have engaged in training, experimentation, showcasing, and responding to criticism throughout the creative activities I have engaged in during my doctoral program, it has been much more than a simple journey towards expanding and honing my techniques. I have consciously chosen to explore a medium—software—which I believe offers a singular opportunity for creative growth, and a worrisome threat to that growth. The emergence of the digital medium as a site of art-making influences and at times disrupts the system of social relations through which people engage with

one another.

I believe that, no matter what roles we have taken in music (e.g. composer, performer, or listener), the digital medium introduces new methods for initiating and evolving interpersonal social relationships, because of the unprecedented capacity of the digital medium to represent and transport animated human agency. This dissertation, and the methods that I have used in my blended role as a composer/performer/programmer/listener, have indicated the potential for changes in the understanding of what music is, how it may be produced, and what it means to exert agency through it. Through a process of research-creation, I encountered challenges to my conception of my own agency over and within music, that formed blockages I identify collectively as a “crisis of agency.” The path to resolving those challenges emerged from my exploration of creativity in algorithmic composition, and the manner in which acts of coding could transform the means by which music could sustain a meaningful interplay of agencies. My methods, drawing on my interdisciplinary background in music composition, programming, the natural sciences, and psychology, helped me in my efforts to extend the form and substance of my music-making, and so engage in the challenging process of letting go of false assumptions of my own role in music, which, once revealed, allowed me to explore the relationships between people musicking through and with machines in a novel manner.

In this chapter, I will describe the interdisciplinary methods used in creating and interpreting my doctoral artworks. I begin with a brief personal biography, to establish the interdisciplinary roots of my practice.

3.2 Origins of my Practice and Perspective on Music

I feel that I emote through music more effectively and honestly than I do through any other medium, artistic or otherwise. While music is abstract, and typically conveys meaning that is not objectively defined¹, I find that abstraction to be a strength. Abstraction gives music an ambiguity, while making composing an act of finding balance, while circling around an undeclared meaning that the listener constructs in their own mind, according to their own musical tastes and enculturation²

1. Consider, for example, how hard it would be to clearly communicate the contents a grocery shopping list through music without lyrics.

. By not directly stating my thoughts when I compose, the process becomes an exercise in sculpting sound with the aim of raising emotion, and the feedback loop into which I enter when working in Cubase is a improvisational cycle of inputting musical material, listening to how it affects me, and then refining, erasing, copying, or starting anew. The process involves a constant modulation of my emotions, and when I am composing well, the computer and I can be seen in cybernetic terms. A cycle of musical note data passes between me and the machine—I input content, the computer plays it back for me, I experience an emotional response, and I input more content. I can lose myself for hours in this cybernetic coupling, for the thoughts and emotions it brings forth in me.

This cycle functions because I feel music deeply, and indeed I experience emotions most profoundly when I compose in my digital audio workstation, Cubase. That said, I also have strong reactions to others' music, and have since childhood. At times I feel that I understand music more deeply than I understand other people. A work of music, at least in the form I conceptualized such artworks prior to beginning my doctoral studies, involved expression and order. Musical grammars served to organize emotions and create narratives that made sense in the absence of semantic expression. The symphony concert was, to me, a high expression of this, where dozens of musicians would perform before a silent audience, all engaged in the evocation of composers' expression as dictated by the score, magnifying social harmony.

My professional biography from prior to my doctoral program reflects my focus on composing, and my willingness to adopt a position set apart from others whose musicking crossed with my own. I am largely self-taught as a composer, and while I studied keyboard performance as a child and am competent with working out my ideas through performance when composing at the computer, I have shied away from sharing my music as a performer. I have had three composing mentors in my life. The first, composer Norman Sherman, was a follower of Schoenberg who composed with tone rows, but who recognized that was not my path, and offered feedback and encouragement throughout my early development, while I was studying biology and then music psychology at Queen's University in Ontario.

Upon graduating with my master's degree in the psychology of music in 1999, I

2. "Musical enculturation" refers to a process whereby an individual absorbs the musical norms of her surrounding culture through both study of music and passive listening.

returned to my hometown of Regina, and over the years developed a practice as a film and television composer. I have also composed concert works that have been performed by the orchestras, chamber ensembles, and choirs³, and have helped to organize and score major arts projects⁴. I am a member of SOCAN, an Associate Composer with the Canadian Music Centre, and I have served on the board of the Canadian New Music Network. I have continued to work sporadically as a composer during my doctoral program, releasing albums⁵.

In my pre-doctoral compositions, I tended to frame my role as being the originator of the score, and the primary source of its meaning. My doctoral program, begun in 2013, gave me the opportunity to test and refine that view, and I produced new, interactive works of algorithmic music, featuring several approaches including those of machine learning (Chapter 5), site-specific installation (Chapter 6), and video games (Chapter 7). These most recent projects, although informed by my practice as a composer of cinematic scores, necessitated a reassessment of my role as a composer. Through changing my medium of self-expression from notation to program code, the behaviour of my agency changed, and I had to locate myself in the new network of social relations formed by interactive, algorithmic music. To do this, and so pass through a nebulous loss of identity, I would have to approach creative engagement as an agent acting both as composer and programmer. In that duality, I sought to form an understanding of my identity that would affirm my choice to locate compositional creativity in the design of algorithms, and find in such acts the same spirit that animated my linear scores, intact and empowered.

3.3 Research-creation Approaches

This doctoral dissertation accounts for the most recent, and possibly most challenging, methodological shift I have undergone in my role as an artist/composer. My exploration of this transformative process uses the creation of artworks as a means

3. These ensembles have included the Regina, Winnipeg, and Victoria Symphony Orchestras, and the Luther Bach Choir

4. These events have included the artist program of the 2005 Canada Summer Games and the “Crossfiring” (2006) event at the Brick Plant in Claybank Saskatchewan, both of which featured dozens of accomplished artists working in diverse media.

5. Two won Western Canadian Music Awards: “Cinematic Symphony and Gamescores” (2013) and “Along the King’s Road” (2016)

to explore ideas of agency, ownership, and expression. I adopt the lens of “arts-based research” to carry out this investigation.

The term “arts-based research” (one of many denoting variations of the same idea) encompasses wide and sometimes controversial understandings of how substantive research may be conducted through creative action. Arts-based research is a diverse conglomeration of research approaches that have gained traction in recent decades as approaches for carrying out qualitative research (Wang et al. 2017, 7). They “emerged as a concept and practice from the interaction between art and social science” (Savin-Baden and Wimpenny 2014, 1) and, since my creative and scholarly work engages with ideas from several distinct disciplines in an interdisciplinary space, I find the integrative aspect of arts-based research appealing.

There are numerous characterizations of arts-based research. While some, such as Patricia Leavy 2009, position arts-based research as a means of gathering data to use in social science or humanities research, my approach is more consistent with Savin-Baden and Wimpenny’s view, defining “arts-related research” as “research that uses the arts, in the broadest sense, to explore, understand, and represent human action and experience” (1). Shaun McNiff’s perspective is consistent with these views. He defines arts-based research as “the systematic use of the artistic process, the actual making of artistic expressions in all of the different forms of the arts, as a primary way of understanding and examining experience by both researchers and the people that they involve in their studies” (McNiff 2008, 29). Since I use my doctoral creative projects as opportunities to explore agency, and the relations among agents, McNiff’s perspective is compatible with my methods documented here.

In more specific terms, I consider my doctoral contributions to be consistent with Chapman and Sawchuck’s conception of “research-creation,” which they describe thus:

Research-creation “theses” or projects typically integrate a creative process, experimental aesthetic component, or an artistic work as an integral part of the study. Topics are selected and investigated that could not be addressed without engaging in some form of creative practice, such as the production of a video, performance, film, sound work, blog, or multimedia text. ((Chapman and Sawchuck 2012, 6))

My work is quite compatible with this definition. I have created a body of experimental works which, taken in chronological sequence, form my attempt to find a balance of agencies within generative, adaptive, and interactive computer music applications that affirms the action of the composer/programmer while granting the interactors/users meaningful interactions with the work among the other agents who act within it.

The four main works of algorithmic music presented in this dissertation are new kinds of experiences through which I challenge dominant narratives such as *Werk-treue* and musicking, and they are expressions of the process through which my understanding of composerly agency was destroyed and rebuilt. I chose not to engage in a scientific study with the aim of achieving empirical results. The evidence of my findings is expressed in the form and function of the works themselves, which, in their chronology, trace the emergence of my conception of the composer's relationship to technology, and how this relationship can benefit human agents who encounter one another through artworks that emerge from it.

As my own doctoral studies have progressed, the products of my creative practice have converged upon systems involving interactive experiences, realized through purposeful design of computer software, leveraging their interactors' appreciation of music to elicit meaningful emotional responses in those who encounter them. Through a process of letting go of the control that, in linear scores of Western Tonal music, is typically ceded to the composer, I have created a series of works that reveal in their form my process of testing configurations of musical agencies. I believe this work is valuable because technological advancements are generally affecting all participants in the experience of music. In particular, as the automation of compositional processes is refined, composers may begin to find themselves in competition with automated systems. The works presented here are the trace of my process of accepting and then embracing automation as a means of magnifying self-expression, and so might offer some hope, and possible ways forward to other worried composers.

3.4 Interdisciplinary Approach

In composing through acts of coding, an individual may draw on skill sets associated with disciplines that lack a long, well-explored history of combined action. The

origin of the discipline of computer science has been located in the 1950s (Denning 2005, 27), and its most fundamental question is “what can be (efficiently) automated?” (Denning et al. 1989, 12) Music, on the other hand, resists simple classification, and has been practiced for millennia by innumerable cultures, with evidence of musical behaviours and technologies (such as bone flutes) extending deep into prehistory (Tuniz et al. 2012; Atema 2014). One of my chief concerns as a doctoral student was to find common ground between the two—techniques that would bridge musical practice and code design to form a practice where I could identify as both simultaneously. This section of my thesis describes how my interdisciplinary background has energized my creative practice, and I aim to justify the methods that I have deployed to develop my art objects and approaches.

My interdisciplinary methods are drawn from perspectives and research strategies associated with the sciences and the musical arts. My practice, however, does not belong entirely to either domain, as I am conducting my research in a personal space that I have defined for myself. As a composer, I have long used digital audio workstation software that models the compositional medium of pen and paper, as used to craft linear concert scores⁶. By this I mean that the representation of the score in the software functions essentially as a written paper score—in basic terms, the user specifies a sequence of note events that, when realized, form a linear composition. There is much added to this process, the Cubase environment offers tools and functions for manipulating note data far beyond those afforded by the simple pencil, but the essence is arguably the same, replicating, as it does, a notated score. I have also programmed since early childhood, and have studied computer science in every degree I have taken. Before my doctoral program, I generally used programming to solve problems, such as coding an application for automatically gaging when a certain type of microbial test kit was first displaying a positive reaction. My education in computer science includes undergraduate-level courses which were intended for computer science majors at Queen’s University, and I have successfully completed graduate computer science courses at the University of Regina in interactive media

6. This is not to say that the workstation acts entirely as a notated score—it is closer in its design to a mixing board integrated with a tape machine playback head. Nevertheless, I maintain that in my own practice, the digital audio workstations I use (particularly Cubase) mirror the notated score to such an extent that I regard many of my actions within such compositional environments as being consistent with those of composers who work in the medium of pencil and score paper.

design and computational audio.

My term project in the former class developed into “Q-Chords” (described in Chapter 5), and I presented that work (Cullimore, Hamilton, and Gerhard 2014) at the joint International Computer Music Conference (ICMC) and Sound and Music Computing (SMC) conference in Athens, Greece in 2014. I also contributed a poster and article to the SMC 2015 conference in Maynooth, Ireland (Cullimore 2015). These computer science conferences add to my presentations at arts themed conferences including a conference of the American Musicological Society held in Lethbridge, Alberta in 2015, and the “Confounding Expectations: Musical Intersections” graduate conference of 2015 at the University of Calgary. I also represented the University of Regina at the 2017 Symposium of the Social Sciences and Humanities Research Council, where I presented my original interactive music systems at a booth. As such, my process includes presenting and publishing in contexts that intersect with both the sciences and the arts. My doctoral studies offered me an opportunity to test that binary distinction between my composing self and my programming self. I began to consciously explore projecting my self-expression through coding, and of analyzing my composing methods to identify the underlying thought processes and techniques that allowed me to compose a concert work or film score.

I framed this convergence of my composing self and my programming self as a conscious integration of two agentic self-concepts that were divided in my mind. As a composer, I felt that I was working with capacities of creativity, inspiration, and self-expression that could not be reduced to formal rules or logically-arrayed sets of processes. On the other hand, I saw programming as essentially logical in nature, involving the application of mathematics to processes for which one has an articulated goal, and conditions under which that goal would be achieved. These opposed perspectives are, of course, overly reductive, and even when I started my doctoral program I recognized that one could be self-expressive through coding and define a logical process for some task in music. But even that understanding segregates composing from programming, and as I progressed through my degree, it became more and more apparent that how I acted as a composer/programmer, and how I understood my nature as an agent within a work of algorithmic music, would have an immense impact upon the quality of my work, and its impact on listeners and interactors who also became agents through encountering it. For this reason, I accept

the dualism of science and art in my mind as a means of analyzing the play of agencies in my work. My journey to finding this integrated self was like a McGuffin—a sought-after object in a film that propels the plot as an idea rather than an actual prop that is used by the characters. In my practice, seeking the self where my inner artist and scientist were joined was part of a larger question, that of asking how I could be an agent in my own works who shared that agency with others in a satisfying, meaningful way. My interdisciplinary identity drove that process of discovery, even as it revealed to me that I was not two agents in conflict, but rather one whose task was to find himself in the musical algorithm, and share that space with others.

3.5 Methods

I will now outline the core methods that I employ, and my main arguments for why I have adopted them.

3.5.1 Composition In Computer-based Environments

I have already described the emergence of my composing practice, from its beginning at the Bontempi organ through my first scores composed in a germinal Commodore 64-based DAW, and on through a gradual evolution to becoming a professional composer. Composition is at the core of my past practice and has been the focus of some of my most significant creative achievements. Experiences where I scored films, produced concert works, and helped to organize arts events helped me to form my identity as a composer. I now see composerly identity in agentic terms, where the composer is an agent whose various capacities (e.g. skills, experience, access to technology, imagination) grant her means to project her agency through production of music. As a reluctant performer, my primary window into musical expression was the digital audio workstation software that I used on my desktop computer. In situating my computer between myself and the audience, my composing tools reinforced my conception of the composer as isolated from audience and performer. There was always technology intervening between myself and the sounding music that I claimed as my own. Indeed, the doctoral projects outlined in this thesis can be understood as the trace of my inner conflict over letting go of this view. It might seem ironic then, that my past use of the computer to compose actually set up conditions that

constrained my attempts to use computers to compose. In the former case, composing was conducted in a digital audio workstation, and in the latter, through coding algorithms. It is the tension that formed as I shifted the location of my compositional creativity from one computer-borne medium to another that obstructed, then energized, my creativity, which I take as evidence of the singular capacity of the digital computer programmer to affect us as agents through choices made in the design of software.

3.5.2 Comprovisation

There are elements of both composition and improvisation in my set of methods. Both of these ways of making seem to be connected. Yet, while descriptions of improvisation may sometimes cast it as “instant” or “spontaneous” composition, Fred Frith notes that the inverse, where composing is defined in terms of its relation to improvisation, is far less commonly observed (Bowers 2002, 10). I draw upon Trevor Wishart’s concept of “slow improvisation” and the idea of “comprovisation,” both outlined in Section 2.7.9, in characterizing my compositional and practice, in both of its forms, linear and dynamic.

When I create art, be it through composition, writing, drawing, or programming, I value both planned and unplanned occurrences. For example, in composing a linear score, even one for a film where I have already confirmed my approach with the director, the software environment in which I begin to work tends to provoke exploration, even when the music I am composing has already been constrained to some goal that the director expects to be fulfilled. The moments that occur between the planned events—such as when I enter the notes of a melody incorrectly, or suddenly imagine an unbidden theme that emerges from the playback of the incomplete cue—are among the moments when I am most creative. I harness processes that I cannot plan, such as the expressions of a distracted mind, inputting mistakes, the fleeting impression of a motif, or the improbable synchrony between playback of the cue and a blaring car stereo passing outside. Occurrences like this are not interruptions, they form points of contact between artist and inspiration, or between conscious articulation and unconscious interpretation. To try to systematize composition in the form of an algorithm might act to deny such inspiration, in the same way that if I tried to compose entirely

according to a plan, the composition would simply recreate itself, contributing nothing new to the world. Comprovisation, then, is an engine of creativity that disrupts my planning and renders the movements of my agency unpredictable in the service of discovering the new. As I completed my doctoral studies, I relinquished control in the service of the music, where before I had attempted to accomplish that very task through exacting control of music through code. My adoption of comprovisation as a composing programmer is thus a key shift in my methods that helped my algorithmic composition to breathe and grow.

3.6 The Computer and the Composer

By deploying the methods I have outlined here, I have evolved my ability to achieve expressive and technical goals. But it is also important to acknowledge that these methods are being carried out in a domain of digital technologies from which many unexpected developments may emerge. During much of my doctoral studies I had increasingly experienced misgivings about integrating computer software and programming practices into my methods of communicating and relating to others through composing. Many have expressed similar concerns, including Hofstadter’s previously-noted discomfort at emotionally reacting to a computationally-generated score (Hofstadter 2001, 4). In the field of artificial-intelligence-produced music, some envision a future where the computer becomes a collaborative partner in music composition in popular music and (in a more innovative application) on social media (Avdeeff, n.d., 10). This process casts composition as a collaborative process between software and human composer, but the consequences of adopting such innovative practices are as yet unknown⁷. What could the consequences of such technologies be?

Writing in 2016, Leo Marx argued that the term “technology” itself should be understood in terms of its effects upon human action, and not merely as a collection of “things” that happens to include machines and computers (Marx 2010, 576). At the same time, Marx observes that technologies are often credited as the source of significant changes through a misattribution of the effects of their use towards the

7. A glimpse of such a future, however, might be afforded by examination of newer generative music systems like Endel, and projects like Autechre’s Max-driven “Confield” album. However, since widespread adoption of generative music systems has not yet occurred, such projects only present a tantalizing glimpse of how music might change when musical algorithms are finally unleashed.

technologies themselves, as opposed to the users of those technologies. He describes this error as involving the attribution of an independent agency to technologies, where there is none. (Marx 2010, 576) He gives as an example a common-sounding historical cliché, saying “the cotton-picking machine transformed the southern agricultural economy and set off the Great Migration of black farm workers to northern cities.” (576-577) Of course, a mechanical cotton picker’s primary concern would actually be to obey the laws of physics; it is otherwise largely disinterested in the social and cultural consequences of its actions for humans. Yet, through misattributions of human agency to technology, Marx argues, we may forget that the impact of machine agencies is granted force by the inertia of the human will, a false but widespread concept that Marx regards as “hazardous.” He says of this fallacy concerning of the nature of technology:

It relieves the citizenry of onerous decision-making obligations and intensifies their gathering sense of political impotence. The popular belief in technology as a—if not the—primary force shaping the future is matched by our increasing reliance on instrumental standards of judgment, and a corresponding neglect of moral and political standards, in making judgments about the direction of society. To expose the hazards embodied in this pivotal concept is a vital responsibility of historians of technology. (577)

Is the digital computer a technology that is just like the cotton picker that Marx uses to illustrate the misattribution of agency? This is a philosophical question with which many have grappled. It is central to the narrative of the singularity, where machines’ capacity to evolve themselves finally exceeds our own (Good 1965; Vinge 1993; Chalmers 2010). It is found in the sometimes-unsettling image of the cyborg (Clynes and Kline 1995; Haraway 2003; Hayles 1999). It underlines debates over the possible upheaval and alienation introduced by digital automation, in music (Cope 1999; Hofstadter 2001) and elsewhere (Turing 1950; Campbell, Hoane Jr, and Hsu 2002). This question can seem all-encompassing and a full account of the nature of the programmable computer in history and culture is likely beyond the scope of a doctoral dissertation. My dissertation explores one corner of this issue: The meaning, for the composer, of choosing to transplant their compositional practice from the medium

of the linear score to that of program code, where human agency is automated with impacts upon human social relationships formed through encountering such music. I examine the choice to embrace programming code as a compositional medium by creating works that reflect different ways of approaching that choice. I then draw conclusions about the nature of composing through algorithmic media based upon what I observe and experience in creating and sharing those works. This knowledge contributes to the much larger phenomenon of humans forming relationships with and through machines. The chronology of my artworks can reveal, in their evolving support of agentic interplay, a path towards more fulfilling musicking facilitated by a composer who also acts as a programmer, and so provide evidence of the potential for automated composition to magnify a composer's agency, as opposed to silencing it.

Media theorist Marshall McLuhan says that humans do not typically understand technologies before we adopt them, with the consequence that media transforms us before we can react or contain such threats. McLuhan does recognize, however, that such understandings may develop over time, and suggests that “the serious artist is the only person able to encounter technology with impunity, just because he [or she] is an expert aware of the changes in sense perception” brought about by exposure to technology (McLuhan 2013, 8-9). To McLuhan, it is our very senses and nervous system that are engaged and extended when we use technology, and this, perhaps counterintuitively, puts us at risk. McLuhan warns that “once we have surrendered our senses and nervous systems to the private manipulation of those who would try to benefit from taking a lease on our eyes and ears and nerves, we don't really have any rights left” (15). Computers may assume just such a purpose today, as I have discussed in Sections 2.6.2 and 2.6.5, where I introduce research on how video games blur the line between human and nonhuman identity. My study of this question is partly motivated by my desire to identify the risks of using technology without knowing it, and a wish to gently nudge fellow algorithmic composers in a direction where humans meaningfully maintain agency over music, even as we hurry to automate its creation.

Chapter 4

“We The Artificial:” As Machines May Think

4.1 Introduction

The work I present in this chapter is entitled “We, The Artificial,” and consists of sixteen songs I composed using my Cubase-based digital audio workstation between 2016 and 2020. See Appendix A for instructions on how to listen to these songs. The form of these songs is based on contemporary pop songs, although I have also incorporated many elements from my practice as a composer of classical concert music and cinematic orchestral film score. These songs are, like most of my compositions, tonal, but diverge from my earlier work in that they have sung lyrics which I also wrote. These lyrics assume the voices of hypothetical artificial beings from a technologically-advanced, but imminent, future Earth. Each song takes as its premise a condition that a future sentient machine might find itself in. “To The Borderlands,” for example, is sung in the voice of a machine travelling to a forbidden zone where it will form a cyborg with a human. In “The Human Race,” on the other hand, the machine laments the petty aims of the humans they observe, from the perspective on an outside observer who life spans many human lifetimes. The settings and artificially intelligent characters featured in the songs belong to many possible futures, and there is no overarching timeline uniting the diverse stories. Rather, each song is set in its own possible future, and the theme that ties all of the songs together is the question of what it means to be human, when technological advances blur the borders that protect the integrity of the human body, and force us to question the

special, privileged status of human identity. This collection of songs is intended to contribute to a prominent, current cultural dialogue over questions of human alienation in the face of machine automation, and, more deeply, of how humans treat those who are different from themselves. I chose to write “We, The Artificial” in the format of a collection of pop songs because its themes are relevant to all people affected by technology today, so I wanted to make the music accessible to a wide audience. Additionally, the technological theme of the lyrics is effectively evoked by the use of synthesizers and styles drawn from electronica, and I feel that in the balance between my electronica, classical concert music, and film music styles, I find an approach that provides the energy, drama, and subtlety that supports the ideas which inspired the album. Appendix A provides a list of song titles, lyrics, and summaries of the ideas that motivated their composition.

In this collection of songs, the artificial being (who I refer to individually as an “Artificial”) also stands in for the human outsider. In each song, the lyrics are sung in the voice of an Artificial who may be misunderstood, constrained, or marginalized. Many of the songs explore possible situations that might occur when humans project their agency through self-aware machines, and show the consequences of considering those that one has power over to be less than human. For example, “Keep The Light Bright” depicts a persecuted Artificial who projects a “hologram from sky to ground” to “defy a history” of persecution and scapegoating of machines.

“We, The Artificial” represents a parallel line of artistic inquiry to my doctoral work in generative and adaptive music. I frame the central theme of this dissertation as involving an exploration of the transformation of human agency when projected through algorithms. I highlight “We, The Artificial” in my dissertation because, over the four years that I have been writing the lyrics and music of these songs, they have served as a canvas where I have worked through my hopes and concerns relating to empowerment of humans through machines, and the empowerment of machines themselves. The lyrics of “We, the Artificial” became my diary, and the songs emerged as part of my internal dialogue about what it can mean to be different, powerful, or ignored. The premise of each song is predicated on a misunderstanding, a disconnect between humans and the Artificial vessels of their agency, that constitutes a harm to the Artificial. As such, this collection of songs is a call to raise empathy. The fostering of empathy is a key aim of my work in generative and adaptive music as

well, particularly the interactive music installation I produced to support an exhibit of Indigenous history and oppression called “Chief Paskwa’s Pictograph,” and my culminating doctoral project “Shards Of Memory.” This chapter, then, acts as a bridge between my established practice as a composer of linear music, and a composer who acts through code, but also as an alternative exploration of empathy, which informs the interactive works.

4.2 Technologies Used to Develop *We, The Artificial*

I composed “We The Artificial” using my Cubase-based digital audio workstation based on a PC computer with 64 GB of RAM, an Intel CoreI7 processor, and Microsoft’s Windows 10 operating system. The instrumental timbres, synthetic and acoustic, were rendered entirely in the computer, using a combination of sample libraries and software synthesizers. Prominently-featured sample libraries included Spitfire Audio’s Chamber Strings library, Native Instruments’ Battery, and Toontrack’s EZDrummer 2. Sampled vocals, particularly sourced from choral and operatic performers, were also featured, drawn from libraries by 8Dio (Insolidus Choir, Studio Sopranos) and East West Sounds (Voices of Opera), among others. These acoustic samples were blended with software synthesizers, including the versatile Omnisphere by Spectrasonics (also capable of synthesis using samples), and Arturia’s collections of emulations of classic synthesizers by Moog, Prophet, Oberheim, and numerous others.

I wrote and performed the sung lyrics myself. I am not a trained singer, but from the earliest stages of the project I felt it important for me to sing the lyrics myself. There were two reasons. First, my workflow with Cubase involves a great deal of experimentation, so the ability to produce, evaluate, and transform lyrics alongside my established composition methods was of benefit. The second reason is that since the album is effectively a diary where the voices of the *Artificials* mirror my own experience, I should perform them myself.

4.2.1 The Choice to Sing as a Machine

The choice to sing my own lyrics, to associate them with my history, may seem contrary to Barthes’ view as presented earlier, where the meaning of the words is a

product of the mind of the reader (or here, the listener), and my presence underlying the textual content of the lyrics does not add or subtract from their meaning. While I agree that the lyrics will ultimately be pieced together in the minds of those who hear them, and my history might thus be a distraction in the coalescing of meaning that occurs in each listener’s mind. I would counter by saying that I am not only writing a text, I am performing it, and that my construction of the music and lyrics alike forms a “slow improvisation.” That term was coined by the composer and software designer Trevor Wishart in recognition of the improvisational quality of composing (Vassilandonakis and Wishart 2009, 10). For me, the computer is essentially an instrument, and I concur with Wishart that my compositional processes resembles an improvisatory performance, emboldened by computer-granted capacities such as “copy-paste” and “undo” that transcend pencil-mounted erasers. As a performance, it is also an act of musicking as understood by Christopher Small, albeit one that reaches its fellow musickers through a recording. In his essay, Barthes says that “writing is the destruction of every voice, of every point of origin” (Barthes 1977, 142). To me, lyrics are not simply words, but the product of an improvisation, where my conflicting beliefs (from “the machine is empowering my self-expression” to “the machine will sap me of my precious creativity”) counter one another in a software environment that introduces the agencies and values of their designers as participants in the creation of the works. Thus, in “We, The Artificial” I am not writing lyrics but performing them in a slow improvisation, and so I disagree with Barthes’ contention that nothing of me exists meaningfully within them.

4.2.2 Creating a Machine Voice of My Own

Initially, I experimented with Vocaloid software as described in the Literature Review, though I did not use the Hatsune Miku library introduced earlier, since I preferred to use libraries designed specifically for English performances. Using the Vocaloid software plugin I was able to craft sometimes-intelligible performances of my lyrics with expression facilitated by automation of plugin parameters. The sung parts projected a robotic quality I felt appropriate for the project. Upon reflection, however, I observed that the Vocaloid singers’ style was baked into the samples, and, at least when I used them, all performances felt like a poorly-sung, glitchy emulation

of a human performance, and I was looking for an expressive, authentic-sounding robotic one.

I began to experiment with vocoder effect plugins. Vocoders were originally invented to reduce the bandwidth required to transmit the human voice (Wilmering et al. 2020, 12), by removing audio information at the point of transmission that could then be resynthesized at the point of reception (Mills 2012, 111). The effect was eventually extended to film and music projects—examples include the synthesized voice of the deranged artificial intelligence “Hal” from the film “2001: A Space Odyssey,” and Wendy Carlos’ vocoder-processed renditions of choral performances from her score to “A Clockwork Orange” (Heesch 2016, 61). In pop music, many artists have used vocoders. The influential German band Kraftwerk introduced vocoder effects in their 1975 track “Pineapple Symphony” (Adelt 2012, 367). Cher’s 1998 hit “Believe” featured a vocal performance processed through the Digitech Talker plugin, which gave its melodies an unprecedented “stepped” quality as well as a tinge of inhumanity (Dickinson, n.d., 334, 336). In my own use of the effect, I initially struggled to achieve the intelligible voice I desired. Eventually, however, I discovered iZotope’s “Vocal Synth 2” software, a vocoding plugin with four distinct synthesis modules. It allowed me to perform my lyrics by recording with my normal speaking voice and to determine melodies subsequently, through feeding MIDI note data into the plugin in sync with the spoken lyrics. This technique meant that the lyrics were discernible from my vocal performances, but the melody, including pitch and timing information, was defined by the midi data, with the result being that I could record a spoken set of lyrics and then later define the melody as a midi track, which was freely editable and subject to my improvisational processes at any subsequent point in composition. Using the Vocal Synth plugin’s “biovox” module, I was able to achieve a sound that, to me, produced intelligible vocal performances that fused human vocal characteristics (such as an adjustable formant frequency and degree of breathiness), with the quantized timings and tinge of synthetic sound that I intended to help the audience distinguish that the voice was almost, but not entirely, human.

As I experimented with vocoder plugins while developing “We The Artificial”, I found that shifting the parameters of my voice allowed me to disappear, my identity replaced by another entity acting as the singer. Unlike the vocaloid, which is exhaustively sampled from the performers of an individual singer, a vocoder relies on the

combination of two audio streams, such as a voice and a synthesizer¹. By carefully selecting and refining various synthesizer sounds to join the stream of my half-spoken, half-sung performances, I could alter the sonic qualities of my voice beyond their natural limits. With a raised formant, for example, I could assume a childlike or feminine aspect that was beyond the scope of my unaffected voice; lowering the format rendered it fuller and grainy. By carefully programming the synthesized audio source, I could likewise brighten or darken the sound beyond the limits of my own voice. Songs the songs of “We The Artificial” are sung in the voices of many machines with diverse designs and experiences, I found that the flexibility of the vocoder permitted me to tailor the expression of the voice to the character whose voice sang each song.

As the vocal technologies and instrumental sound palette became more established, I continued to compose music and lyrics. Songs were written in the Cubase digital audio workstation, where I input and edited events in environment’s “piano roll” editor. Aside from the vocal performances, which were a new development in my musical practice, the essential tools and techniques mirrored or elaborated on those I had used since my earlier compositions with the “Ballade” sequencer I first used in the early 1990s.

4.3 Motivations For “We, The Artificial”

I have already pointed out that electronic technologies, including the digital computer, have become an important site for making and sharing music. Messiaen’s use of the Ondes Martenot in his symphonic scores, Wendy Carlos’ arrangements of J.S. Bach’s music for the Moog synthesizer, and the relocation of the composition of film scores from the traditional medium of the handwritten paper score into the computer-borne domain of digital audio workstations help to form a larger bridge through which human self-expression and technology are united. “We, The Artificial” extends that premise, by exploring human self-expression as it might appear if it originated from machines themselves.

1. In this process, the stereo recording of my sung lyrics, called the modulator, is analyzed by the vocoder plugin during playback to create a series of frequency-specific volume envelopes which then are applied to the carrier, a synthesized tone. The carrier signal is broken into frequency bands, each of which is subject to modulation by the envelopes of corresponding frequencies derived from the modulator. A helpful introduction to the vocoder’s function and use is provided here: <https://www.musicradar.com/how-to/how-to-get-the-most-out-of-vocoders-a-complete-guide>

I chose this premise partly because I feel like I share qualities often associated with machines. I appreciate the beauty of mathematics and the order of numbers, and find the act of counting relaxing. I sometimes struggle to understand others' emotional responses, which recalls the trope of the unfeeling artificial intelligence (e.g. the soul-searching character Data and the soulless Borg hive mind in "Star Trek," as well as the clinically genocidal Skynet and its murderous android minions in the "Terminator" franchise). In my darker moments, I sometimes feel a disconnect with others, as if my thoughts are somewhat out-of-sync with those of others. I respect the machine intelligence because, though it is the product of humans, it can transcend us, surprise us, and teach us. The sliver of agency we share with them becomes autonomous, it grows apart from us and becomes its own phenomenon. Machine intelligences may well outlive us, and however that split occurs, through the gradual fading of biological humanity, or a cataclysmic war of annihilation, it is something that I am in awe of. Whatever happens, however, the machines will ultimately be a product of human agency, and we shall live on as long as they do.

4.4 Influences on "We, The Artificial"

I see two areas where "We, The Artificial" innovates with respect to music. One is in its particular configuration of musical influences and form. The other is in its vision of the artificial being.

4.4.1 "We, The Artificial" as a Concept Album

The composition of "We, The Artificial" took me to a space between genres at a point where art music and popular music encounter and change one another, consistent with the popular music form known as the concept album. A "concept album," as defined by Lori Burns, is "an album that sustains a central message or advances the narrative of a subject through the intersections of lyrical, musical, and visual content" (Burns 2016, 95). Examples of concept albums span decades of popular music, with early examples including the Beach Boy's "Pet Sounds" (1966) and the Beatles' "Sgt. Pepper" (1967). The form was highly popular during the era of art rock and progressive rock, which saw the release of significant albums such as David Bowie's "The Rise and Fall of Ziggy Stardust and the Spiders from Mars" (1972) and Pink

Floyd's "The Wall" (1979). In more recent years contributions such as Taylor Swift's "1989" (2014) have introduced the form to contemporary youth audiences (Burns 2016, 94).

Some see the concept album as a bridge between popular music and art music. Katrina Gingrich, for example, positions the concept album in the broader classical tradition of the song cycle: "A concept album possesses all the defining characteristics of a song cycle but expresses them in more modern language: three or more defined sections that use text painting and a logical sequence of tonicities to express a single poetic theme" (Gingrich 2010, 71). She notes, however, that the concept album diverges from the song cycle in two ways: that that the concept album is primarily a recorded medium, where the song cycle may be performed, and that classical song cycles usually have only one voice or narrator, whereas a concept album may have many.

The Who's rock opera "Tommy" (1969), and Kraftwerk's concept album "Computer World" (1981) have been key influences on my composition of "We, The Artificial," so a brief description of them is warranted.

4.4.2 Influences: "Tommy" and "Computer World"

From its inception, "Tommy" was lauded as transcending rock music: in an 1969 article in "New Music Express" published prior to its release, the double album was anticipated as "probably the most important milestone in pop since Beatlemania. For the first time, a rock group has come up with a full-length cohesive work that could be compared to the classics" (Sanders and Dalton 1969). "Tommy" began as a concept album composed by the Who's Pete Townshend , and following its early success it was adapted into a 1974 film by Ken Russell, adapted to a setting in the aftermath of World War II, and featuring a cast of film and music celebrities who, the producers felt, would draw in wide (and overlapping) audiences of youth, people sympathetic to establishment and anti-establishment politics, and African-Americans (Leverence 1974, 470).

The film adaptation of the original concept album became something of an obsession of mine while I was first learning to compose as a teenager. I watched my VHS copy countless times, and I believe that my viewing to the film, in which the

original songs from 1969 were arranged for a blend of orchestra and rock instrumentation, helped to nurture a core method in my approach to composition, the fusion of popular and art music elements. More significantly, some of the central themes of “Tommy” are echoed in “We, The Artificial.” In an analysis of the original concept album, Graeme Daniels exposes the psychological subtext of “Tommy.” Tommy was famously rendered “deaf, dumb, and blind” by the trauma of witnessing of a murder that the mother and stepfather demanded he not speak of or remember. Tommy, according to Graeme Daniels, is actually dissociative, unable to reconcile the memory with the lie he was forced to carry (Daniels 2016, 101). Tommy’s retreat inward, according to Daniels, forms “a haven, but ultimately an unsatisfying, perhaps terrifying state of self” (102). In “We, The Artificial,” the machines similarly stand in for people dialoguing with altered mental states. In “Singing Of Integers,” for example, the voice is that of an Artificial who finds comfort in a compulsion to count, a tendency that allows that Artificial to represent those with autism or obsessive-compulsive disorder, conditions which often are associated with repetitive behaviour (Zandt, Prior, and Kyrios 2007, 251). Other Artificial characters echo this detachment from reality, since they are not only outsiders, but also the product of a science that unknowingly harms them. The character of Tommy shares a core aspect with many of the Artificials: an experience of being different from others that propels a drive for interpersonal connection.

Unlike “Tommy,” where an inward-looking individual finds salvation in connecting with others, “Computer World” (1981) depicts a world where humans become more like machines. “Computer World” is a concept album by the innovative German band Kraftwerk. Kraftwerk has been described as a “conceptual band,” or “Gesamtkunstwerk.” The term was popularized by Richard Wagner in the Nineteenth Century, where it came to mean the unification of all art forms within theatre (Groys 2018, 69-70). Applied to Kraftwerk, it suggests that the image of the band, in all its forms (e.g. music, album art, concert staging, and promotional material) was uniform (Adelt 2012, 368). For Kraftwerk, that image assumed the identity of the robot, which was partly a reference to the then-prevalent stereotype of Germans as “robotic automatons” (371), but also a reflection the band’s interest in machines. Kraftwerk celebrated machines in their references to technologies such as cars (“Autobahn”), spacecraft (“Spacelab”) and, later, computers (“Pocket Calculator”). Their album

“Computer World” particularly resonates with “We, The Artificial.” Its vocoded vocals evoke an “Orwellian” world where surveillance is rife (in the title track) and humans seek comfort with robotic lovers (in the song “Computer Love”) (Adelt 2012, 371). The lyrics of “Computer World” are often simple and repetitive, but rather than detracting from the music, I feel it reinforces the machine aesthetic. The song “Numbers,” for example, consists entirely of human and synthetic voices counting to eight in several languages, evoking the simple inevitability of the unfolding of an algorithm, and serving as an inspiration for my own song “Singing of Integers.”

What excited me about “Computer World” when I discovered it in the late 1990s was how the hypnotically-repetitive themes evoked a dystopian future where the border between human and machine is blurred and ignored. As I was at the time composing music entirely in the computer-based medium of the digital audio workstation Cubase, I was realizing that the computer was, in a sense, enabling me to express myself, and thus granting me a voice. Like “Computer World,” “We, The Artificial” concerns a state of being charted by Donna Haraway, who observed that “we are cyborgs” (Haraway 2003, 292). She argued that the border which we use to distinguish our human form from the machine is leaky, allowing machines to appear “lively,” and humans “frighteningly inert” (294). Haraway goes on to argue for “pleasure in the confusion of boundaries and responsibility in their construction” (292). In “We, The Artificial,” the artificial being is essentially human, and their alienation is that of the individual who is understood as an object rather than as a being. “Computer World,” by directing its gaze towards an uncertain future where digital technology has a complex and worrisome relationship with humanity, helped me to question the relationships of humans and technology, and how it might us into technology, and technology into us.

4.4.3 Other Influences

“We, The Artificial” recalls many other works that explore the relationships formed between humans and machines, from Fritz Lang’s “Metropolis” (1926) through the long-awaited sequel “Blade Runner 2049” (2017). Other explorations of the nature relationships between humans and machines that have influenced me include the films “Alien” (1979) and “Tron” (1982), Isaac Asimov’s “Foundation” novels, the

utopian, machine-controlled “Culture” featured in Iain M. Banks’ novels, and video games such as “Mass Effect” series by BioWare and Bethesda’s “Fallout” series. This is by no means a comprehensive list, which I feel illustrates how important a dialogue on relationships between humans and machines has become, among artists and in popular culture alike. My first encounter with seemingly-sentient machines, one which I believe introduced many people of my generation to a blended human/machine culture was watching the original “Star Wars” film when it was first released in theatres in 1977. The “droid” characters clearly possessed personalities (e.g. C3-P0’s timidity), and were capable of heroism (e.g. R2-D2’s efforts in the battle against the Death Star), but were still bought and sold like slaves. The characters C3-P0 and R2-D2 resonate with me because they participate in and contribute to their fictional society, but still are not generally granted the full status of the individual. These characters, and others like Blade Runner’s replicants and James Cameron’s Terminators, form the template of the Artificial, as the machine who operates in society but not as a full member of it.

4.5 Reflection: Agency in “We, The Artificial”

This album falls from one central premise: that the self-aware Artificial being symbolizes the human outsider, and that by understanding and accepting the Artificial, we magnify our empathy for other humans. I draw upon this premise most clearly in the song “User, User.” The lyrics are sung in the voice of an Artificial who, recognizing that its wellbeing is subject to the mercy of a user unaware of its humanity, asks that user simply to “be kind to me.” This song is very personal, because it echoes my own fear of losing my agency due to the action of others who might not act with kindness. In a broader sense, however, the song reflects my belief that all beings should be treated with kindness and respect, whether they are human, animal, or machine.

In this light, what might it mean for a human to interact with machine whose behaviour is governed by algorithms? This is the foundational question propelling the work I document in this dissertation. I will address this question in this reflection in terms of how I engaged with it through my process of creating “We, The Artificial.” I will explore three topics in this reflection. The first links the Artificial to the video

game character, both of which can be considered nonhuman participants in human social interactions. The second discusses what I discovered about machine-borne agencies through writing songs in the voice of hypothetical, sentient machines. The third explores how I project my agency as a composer who works in a digital audio workstation environment, and sets up my characterization of how a composer's agency may be transformed when she relocates the site of her creative expression from a notation-based environment to a coding-based one.

4.5.1 A Depiction of Agency

The Artificials exist in society, but demand more than being perceived in social terms. Recall Heider and Simmel's study of human interpretations of simple animated shapes in Section 2.6.2. Even the participants in their study were not fooled into believing that the geometric shapes were actual people, but rather used a model based on social norms to simplify the explanation of otherwise arbitrary movements (Heider and Simmel 1944, 256). In contrast, the Artificials are not merely vessels for agency, they are beings-for-themselves, conscious and emotionally complex, human in all ways except for the materials from which they are assembled. They are not simulated humans, they are humans essentially, and to say that they are anthropomorphized would deny them that distinguishing element of their identity.

The Artificial is not a model for a musicking machine such as the systems I produced that will be introduced in the next three chapters. Rather, it is a model for the composer, who seeks to reach out to others in a spirit of fostering empathy and strengthened social connections. The Artificials' mechanical bodies separate them from humans, both in terms of their capacity to interact, and in terms of how such nonhuman interfaces engage human assumptions of the abilities and worth of the Artificial. The underlying motivation for writing these songs was actually my desire to understand my own agency and how its action might be impacted by my use of technology, while exploring the authenticity of my own creativity when I compose with and through algorithms.

That is the main reason why I chose to include "We, The Artificial" in my dissertation. As I explored the power and limits of human agency, through reading, coding, songwriting, and maturing during my doctoral studies, I came to believe that

agency is precious, and when one is prevented from acting in accordance with their will, through alienation, suppression, or being deemed as less than what we are, it may be a harm. Writing “We, The Artificial” allowed me to explore the bounds of my own agency, and extend them, because in writing the songs I was able to uncover key themes in my own life that had inhibited me. “Miracles And Revolution,” for instance, allowed me to dialogue with my fear of my agency being rendered redundant or diminished by technologies that exceed my composing capabilities in some manner (e.g. by being more prolific, or more engaging to audiences). I dialogued with this fear through writing the song “Miracles And Revolution,” which contains the lines:

These are the times
Of miracles and revolution,
Of upheaval and transition.
As we remake the world,
We know we owe a debt to you—
You are the template for our higher selves,
Thanks for sharing the things you value.

Through these lines, I articulated a resolution to a problem that distressed me through much of my doctoral program: fear of the “singularity,” Vernor Vinge’s apocalyptic point in the near-future when machines can construct and extend their capabilities independently of humans, creating an uncontrollable increase in technological complexity that would overwhelm humanity before causing its extinction (Vinge 1993). Through writing songs like “Miracles And Revolution,” I questioned assumptions underlying my apocalyptic fears. Perhaps we might instead form a continuity with super-intelligent machines, where the more we share of our agency with them, the more of us they carry forward. Articulating this possibility was a comfort, and was part of my process of reframing my doctoral coding projects as an extension of myself, rather than as a reductive systematization of my composing process, which, earlier in my program, seemed an existential threat to my self-expression.

I feel that most human-designed algorithms carry not only the form of those humans’ idealized solution to a problem. They may also carry the values, biases, knowledge, and misunderstandings, of their coders. This understanding informs every song in “We, The Artificial”—that the relationship between the human being and

the function of the machine is a continuity. This idea not only informed the ideas I articulate in the album, it also can be uncovered in the processes by which I produced the album, a topic to which I will now turn as I conclude this chapter.

4.5.2 An Enactment of Agency

I wish to situate this album, and my practice of linear score composition more broadly, more clearly in the creative path I describe in the next three chapters. Here, “We, The Artificial” represents the established methods I have used to compose, from which my coding projects diverge. The form of this divergence, which I characterize as transformation of composers’ agency, is the central topic in this dissertation. Thus a characterization of how I acted to project my agency while creating “We, The Artificial” is warranted.

The Digital Audio Workstation

Central to my composing process is the technology of the digital audio workstation. These audio production environments began as simple, specialized hardware or software audio editors, but have since become key tools used by the modern composer (Marrington 2017, 77). I use the software environment Cubase, which I liken to a word processor, except the interface replicates a sheet of score paper, rather than a blank page for recording text (Cubase is actually capable of much more than MIDI editing, including mixing audio and enabling remote collaboration, but the MIDI engine here is my chief concern). Like a word processor, Cubase does not directly compose for me, but gives me a set of tools I might need to record and organize sonic events in a composition. I have used Cubase for over twenty years, and its interface now feels transparent, offering little resistance to the flow of my work. Indeed, at times Cubase may seem to be a simple tool akin to a hammer, in that it offers me the ability to transparently record and edit my ideas, facilitating rather than directing my actions. I use the example of a hammer because it is inert, extending its user’s arm without controlling where that arm swings, and when I am composing effectively in Cubase, I often feel similarly—that there is no distance between myself and the score, and that my agency is transmitted whole.

This sense of control, however, is something of an illusion. Adam Patrick Bell

argues that Cubase and other digital audio workstations are not truly transparent, as he illustrates with the following thought experiment (Bell 2015, 47). Bell asks us to imagine a teacher distributing a sheet of paper and a pencil to each of her students, then asking them to “make music.” Bell suggests that while some would begin to draw staff lines and pepper their sheets with penciled notes, others might drum upon their desks with their pencils, or explore the sonic possibilities of crumpling or ripping paper². The form of the students’ music is thus enabled by the specific tools they have been given, and were the students given other tools to make music, such as Cubase or a tape recorder, their music would necessarily take alternative forms. This hypothetical scenario illustrates how artists’ tools shape artists’ expressions, and might lead one to question if the transparency of the simple hammer even exists among the many tools which people use to make and share music.

Affordances

Bell is demonstrating that the actions composers take are constrained by their environment, including those objects we deploy in our music-making. He references James J. Gibson’s theory of “affordances,” a useful model for understanding how the composer’s behaviour might be affected by the devices they use. Gibson defines affordances as features that an environment offers an animal, whether they be helpful or hindering (Gibson 1977). When humans change the environment, through reshaping the land or introducing technologies, the effect can be framed as an altering of available affordances (Gibson 1977). Perhaps food might be made more plentiful through agriculture, or a settlement made safer from attack through wall-building. Such environment-changing affordances can actually alter the course of human history—agriculture may have aided the development of cities, and placing walls around them might have supported the apparent cultural divide between rural and urban people.

2. these examples of creative thinking regarding the pencil in music-making contexts point to the value of divergent thinking in digital audio workstations. Arpeggiators, loopers, and other interface components that automate and externalize compositionally-creative behaviours might be seen as limiting where their functions substitute for the application of the composer’s agency over musical expression. That said, such automating features, when seen in unconventional ways, might still offer intriguing expressive possibilities. One might apply an arpeggiator to non-pitch-based variables, or use a graphic editor to manipulate spectral representations of a sound. In such cases, automation is directed into unconventional modes of processing, and so afford the composer opportunities for creative expression that mitigate the loss of agency over musical form that conventional approaches to automation might not yield.

Given that the composer's environment is also full of affordances, software-based and otherwise, Bell chooses to examine the digital audio workstation through Gibson's lens.

In his analysis, Bell observes that digital music technologies are not only widespread³, but that, as affordances, they have facilitated broader public engagement with basic composing practices. Indeed, there are more composers alive today than not, and many of us need not rely on the traditional composer's tool, a writing implement like a quill or pencil (Bell 2015, 46). Yet wide access to digital audio workstations has not necessarily been a boon to music listeners or composers. Thomas Regelski observes that the use of composition software can still be limited by the "musical skill and knowledge of its user," (Regelski 2007, 36) although this group of users represents only a portion of the whole population of composers. Bell concludes from this that technological aptitude is not enough to permit many of those who use Cubase and other composing software the ability to excel at the task of composing itself (Bell 2015, 46). The affordances of technologies like Cubase affect people differentially, in that they may provide many options for action, but the limitations of the user's ability or capacity for expression may render a portion of these potentials inaccessible.

Algorithmic Affordances

When I use Cubase, I am aware that my compositional development is both enabled and constrained as a consequence of the design of the software. Its affordances may shift with each update, as features are added, taken away, elaborated, or simplified. These changes they shape who I am becoming as a composer, and the nature of the music I write, as they alter the available pathways by which I might exert my agency. Bell quotes Mark Ellis, also known as Flood, a producer and sound engineer who has worked with U2 and the Smashing Pumpkins. Flood provides a view of technologies for music creation that I find compelling:

Essentially, I believe that modern technology is an instrument in the music. I am an advocate for technology as long as it's something that's used rather than something that somebody is used by. It's very easy to

3. For example, Apple's *GarageBand* software is available for free on all current iOS- and MacOS-based computers.

become a slave to technology and do something over and over again. Even though you have the ability to try a lot of options with a computer, that doesn't mean that it will be done any better or quicker. It is important that you make sure you use technology to your advantage and don't ever let yourself become used by it.

(Flood, in Bell 2015, 52)

Flood is warning us about creative technologies' ability to change us, the idea that technologies can facilitate human creation, but that this may come at the expense of expression, in that technology can limit us by extending our ability to produce beyond our natural capacity. When I use Cubase, it affords me increased memory capacity, by recording every note I enter and presenting them to me, graphically and sonically, when my recollection of them fails. It affords me an aid to my aural imagery, since it can render my score as an audible performance, negating my need to follow a score visually and infer its sound structure. In some cases, the workflow guides my expressive choices. I might experiment with a variety of timbres for a theme by looping its playback while loading different instrument presets, and, should one appear that I find appealing but didn't anticipate, I may still use it, causing me to diverge from the general path formed from my earlier intent for the composition. While this capacity of the technologies behind Cubase—namely, the constraint and/or redirection of the agency of the composer through the introduction of other agencies into her process—might be seen in the pencil as well, I would draw a contrast. Fundamentally, the pencil is inert. Cubase is not. By this, I mean that a pencil cannot move itself, whereas Cubase can, in that it can mirror aspects of the execution of human agency. It may suggest or enforce decisions on behalf of the user, through an implementation of templates or preset plugin configurations, for example. It might also extend the user's memory at the cost of her mental representation of the score, where a composer is not concerned with the names and relationships between chords and keys in the score, since that information can be heard on playback and used in composition guided my intuition rather than music theory. I would argue that such affordances are not provided by the pencil, because they arise not from the structure of Cubase, but from its expression as processes. Calling up a template, playing back a work-in-progress, and selecting a preset sample or loop all require the computer to act, and action towards a goal requires intent. If this intent does not

come from the user, it must come from the code itself, which can only mean that the intent is a trace of the system, plugin, or template designer, since the code is not able to act outside of their specifications. The pencil conveys no outside agency because it is not enacting a process independently of its user. Cubase, however, does, and that aspect of the system elevates it above a tool, into a site for an interplay of agencies.

Bell cautions against making the “facile fallacy,” which he defines as “the misconception that music can be made seemingly instantaneously and effortlessly with technology” (Bell 2015, 52). He notes that Apple has marketed their consumer-grade digital audio workstation GarageBand with the slogan “Incredible music. In the key of easy” (52). Software designed for ease-of-use, helping inexperienced composers to create music without needing the experience of training or self-directed study, might provide an incentive to such potential composers in the form of a quick path to producing music. But Bell observes that “there is a layer of music education woven into the code of the software that makes the ‘ready-made’ experience possible.” (53) Depending on its design, a digital audio workstation like GarageBand or Cubase might allow its users a rapid path to musical results, but in the absence of experience with aspects of musical education that are displaced to the software, an opportunity for development of expression might be lost. In such instances, the computer may become something resembling a crutch—it helps the user to travel, but at times they may travel less freely, such as when a jump is required. And sometimes, in music as in walking, it is the jumps that take us to unexpected discoveries.

Agency And The Digital Audio Workstation

I feel a tension when I consider how I benefit from the way Cubase extends the scope of my compositional ability, because I have become aware of the power of software to limit or redirect creative expression. The digital audio workstation contains within it what Bell calls a “layer of music education” that may afford its user quicker, more complex, and/or more engaging compositional results (53). In such a case the expertise of the designers of Cubase or GarageBand becomes partly responsible for the form of the music created through these systems. Their agency is, arguably, ever-present in its melding with users’ compositional agency, insofar as it shapes the product of a composer using the system. Digital audio workstations may thus be framed as an affordance that the composer becomes dependent upon to

function fully, at the price of admitting the agencies of the software’s designers into the determination of the form of the work.

I have said that the vocoder plugin gave me my voice in “We, The Artificial.” By using one, however, I became subject to its constraints. For example I could only use one of five available vocoder modules, each of which had a particular, fixed aesthetic quality, which effectively constrained the feel of machine voices I might synthesize. The previous version of that plugin had only four, omitting the novel “Biovox” module that was featured on most of the more recently-composed songs. I chose the Biovox module because of its more “human” quality, and I believe that its increased intelligibility and less mechanical sound helped me to find myself in the lyrics when I played them back. If this made me a bit more honest in my characterizations of the Artificials’ as standing in for my own, I can only wonder what would have happened had the designers of the plugin chosen to develop that module first, in which case I might have found myself in the voice earlier in the project. That choice was not mine, however, it was an artifact of the plugin designers’ agency, and illustrates how a composer’s agency meets those of coders and technologists in their most personal acts of self-expression.

4.6 Conclusion

“We, The Artificial” represents my practice as it existed prior to my experiments with composing music through coding. I characterize my established composing practice, involving the creation of linear scores within the environment of the digital audio workstation, as one where I make use of externalized tools that enhance my composing capacities, at the cost of loss of control. In other words, my expression becomes the product of additional agencies that are not my own, raising questions regarding the integrity of my expression. The theme of the concept album “We, The Artificial” explores this issue by portraying artificial, algorithmically-controlled beings as essentially human, questioning in the lyrics what the distinction between humans and machines might actually be (if any). I reveal through the lyrics my emerging belief that, in theory, humans and machines are functionally similar, and that, since we share a social space in which we implicitly assume that machines have human characteristics, these machines deserve to be treated humanely. The agency projected by

the machine should not be ignored or dismissed, as it contributes to and extends our means of self-expression, through introducing us to others via mingling of agencies. When we act with compassion and empathy towards machines, and recognize their inherent humanity, we avoid guilt and enter into a new, potentially empowering state of being and production. To achieve this, however, we must understand what the machine is, and develop them wisely. If we do not, they may be the agents of our annihilation.

In the following three chapters, I explore these themes as they are raised in the context of my code-based compositional endeavours.

Chapter 5

Q-Chords: Hoarding Agency

5.1 Introduction

This chapter describes my first substantial attempt at a code-based automated composition engine. The program, called Q-Chords, was designed in Max-MSP as a term project for a graduate course in interactive media and video game design. This work was also presented at the 2014 joint International Computer Music Conference (ICMC) and Sound and Music Computing (SMC) conference in Athens, Greece in 2014, and published¹ in the conference proceedings. I wanted to create an application where code would compose autonomously, and, given the topic of the course, I chose to develop a system that could be useful when integrated into a video game.

I considered the kinds of automation problems I would encounter were I attempting to realize a video game score. I began to consider the connective material between cues: bumpers, stingers, and transitional fragments that link more developed themes. Composing such transitional sequences seemed to me as a somewhat more formulaic task than writing main themes. The constraints on the composer of a transition are clearer, given the basic goal of such a sequence and the fact that the form of a transition is partly determined by the form of the cues which it bridges (with respect to instrumentation, key, and tempo, for example). Accordingly, I felt that my first attempt to create an automated composer might be more successful if I began by emulating a task for which I could articulate a clear problem and criteria for its solution.

1. Citation: Cullimore, Jason, Howard Hamilton, and David Gerhard. 2014. “Directed Transitional Composition for Gaming and Adaptive Music Using Q-Learning.” ICMC—SMC—2014: 332-338.

I focussed my project around one compositional task familiar to film and video game composer, that of changing from one key to another. A few elementary concepts of music theory are relevant to this task, so please review Appendix B for definitions of some important terms used in this chapter.

I reasoned that, if a video game composer were to produce a library of cues, they would likely prefer the freedom to write in any key, as I do. Restricting the variety of keys deployed in a video game score is not an uncommon strategy, however. For example, consider the case of parallel composing, where the composer writes multiple versions of a cue that can be seamlessly faded from one to another and, as a consequence, need to occupy the same key (Van Geelen 2017). Tools that could automate transition generation in a manner that allows video game composers to score cues without limiting harmonic centres would be liberating. As such, I decided that my project would involve automated creation of chord progressions that would link any two major chords derived from the keys of two music cues that are to be connected. The initial chord would be set to the final chord of the cue that is ending, and the terminal chord of the progression would be set to the first chord of the next cue. The process of changing keys, or “modulation,”² and, in Western Tonal music composition, modulation is typically carried out with the use of pivot chords and notes. For example, to modulate from C major to *B* \flat major, a composer of tonal music might choose to pivot using the chord F major as a bridge between the two keys, since F major is a common chord in both of the scales of the bridged keys. Over a span of eight sequentially-presented chords, my program would automatically construct an ordered sequence of chords intended to modulate between that last major chord of one cue and the first major chord of the next. I chose to require generated transitional sequences to begin and end with major chords to simplify the development process given the time constraints of the course project and term.

My other goal in designing Q-Chords was to make the approach as general as possible, as opposed to being coded with such specificity that attempts to adapt the code to wider applications—for example, by removing the restriction on the mode of initial and terminal chords—would be incompatible. I had discovered the field of reinforcement learning through my readings for the interactive media course, and felt that an approach where the program could learn to create a desirable harmonic

2. Please see Appendix B for definitions of some of the music theory terms used here.

sequence semi-autonomously would potentially make it more flexible. I also encountered a paper describing composition-generation via a form of reinforcement learning known as “Q-Learning” (Reese 2011), and realized that I could improve upon their approach and adapt it to my own purposes. As such, I began to develop a chord-pattern generator based upon the established Q-Learning framework.

5.2 Overview of Q-Learning

In a discussion of the Q-Learning framework, Pierre Yves Glorennec describes reinforcement learning as concerning “a family of problems in which an agent evolves while analyzing consequences of its actions, thanks to a simple scalar signal (the reinforcement) given out by the environment.” (Glorennec 2000, 17) In other words, a Q-Learning implementation implements an “agent” which acts by travelling through a code-based representation of a space, and which experiences varying degrees of reward or punishment based upon the path it takes within that space. The ultimate challenge for the agent in such a system is “to understand how a current action will affect future rewards.” (Even-Dar and Mansour 2003, 1) Through the application of reinforcement learning techniques, a computer program can acquire knowledge of a domain by traversing it and experiencing the ensuing consequences. I can then use this library of acquired knowledge to solve the pathfinding problem, the question as to which path optimizes the rewards the agent experiences during a single traversal of the domain from its beginning to its end.

5.2.1 A Brief Introduction to Reinforcement Learning

Reinforcement learning is a subset of a broad field of computer science research known as “machine learning”, which involves developing techniques for computers to learn about data and use that knowledge to make inferences about it. (Reese 2011, 33) Reinforcement learning (and its Q-learning variant) are modelled upon natural behavioural strategies observed in animals. An animal may be forced to make decisions that affect its survival or reproductive success, and in order to maximize the possibilities of favourable outcomes, the animal may choose behaviours based on its past experience. (Gasser 2009) An animal may have incomplete knowledge of its environment, however, and as a consequence it cannot utterly know or reliably

predict all possible outcomes before it encounters them. If the animal has sufficient cognitive capacity, it may experiment by attempting a selection of outcomes, facing the consequences of its choices, and then applying what it has learned when it needs to determine its behaviour in subsequent similar situations. Reinforcement learning systems can take a similar approach. The program may form a domain in which the agent is presented with choices regarding where it may travel within that domain. In varieties of reinforcement learning such as Q-Learning, the agent “is not told which actions to take... but instead must discover which actions yield the most reward by trying them.” (Sutton, Barto, et al. 1998, 1.1) The agent may begin by knowing nothing of the domain, other than the subset of states accessible from the current one, into which it may travel. It learns by acting within the domain and developing its understanding of the desirability of various pathways through recording punishments or rewards it receives after making those choices.

Agents act in accordance with a “policy,” which defines the manner in which an agent behaves as it traverses a domain. A policy answers the pivotal question as to where the agent should go next, under all possible contingencies. (Sutton, Barto, et al. 1998, 1.3; Poole and Mackworth 2010, 390) The agent’s ultimate goal is typically to uncover an optimal policy that maximizes the benefit of the particular sequence of choices that solve the pathfinding problem, charting a path from an origin state to a different destination state. (Glorennec 2000, 20) The benefit the agent receives for following a certain policy takes the form of a numerical “reward”, the value of which is proportionate to the desirability of the action the agent has taken. (Sutton, Barto, et al. 1998, 3.2) An optimal policy maximizes the overall reward value of the actions taken by an agent as they traverse the domain. (Sutton, Barto, et al. 1998, 3.1; Littman 1994, 2)

Machine learning frameworks fall into two broad categories: “on-policy” and “off-policy”. (Sutton, Barto, et al. 1998, 5.6) In the case of on-policy approaches, the agent may “evaluate or improve the policy that is used to make decisions.” (5.4) Here the agent acts in the domain with an awareness of what it has learned, and it refines this knowledge through engaging in new actions for which the learned policy acts as a guide. In “off-policy” varieties of reinforcement learning, learning processes occur separately from the determination of agents’ actions, and as a consequence, the optimal pathway uncovered by the agent is not dependent on the policies governing its

behaviour in the domain. (Poole and Mackworth 2010, 470) However, even though the agent’s behaviour may be determined randomly, an optimal pathway can still emerge. The gradual formation of an optimal policy is a process known as “convergence.” This phenomenon occurs over numerous traversals of the domain by the agent, and refers to the gradual decrease in the rate of change of learning data as the values approach the theoretical optimal values. (Lytle 2015)

5.2.2 Q-Learning

The Q-Learning framework was first described by Christopher Watkins in 1989, and further refined by him and Peter Dayan in 1992. (Watkins 1989; Watkins and Dayan 1992) Q-Learning implementations may be understood as “off-policy” approaches to reinforcement learning. (Poole and Mackworth 2010, 470; Sutton, Barto, et al. 1998, 6.5; Even-Dar and Mansour 2003, 1) This attribution is valid because an agent in a Q-Learning implementation requires no knowledge of the value of actions within the domain, either predetermined or acquired at runtime, to converge upon an optimal policy for traversing a domain.

At runtime, the agent engages in a series of traversals of the domain, beginning in a specific origin state and ending in a specific terminal state, the entry of which concludes the agent’s travels. See Figure 5.1 for an illustration of a simple six-state domain. A complete traversal of the domain from origin to terminus is known as an “episode.” (Sutton, Barto, et al. 1998, 3.3) During its travels, the agent may only occupy one state at a time, but may pass from state to state by taking an “action”. (1.1) The states into which the agent may travel during a given action may be restricted to a subset of the entire domain, with some states available for entry from the current state while others are inaccessible. The agent traverses the domain through a process known as an “experience”. An experience has three components, state-action-reward, (Poole and Mackworth 2010, 468) and unfolds as follows:

- *State*: The current state in the domain is identified, along with the set of permissible destination states accessible from the current state.
- *Action*: The agent chooses a new state to enter from the set of possible states, and transitions to the new state.




Figure 5.1: A domain of six states. Arrows indicate state-to-state transitions available to the agent.

- *Reward*: The agent receives a numerical reward, which may be low (punishment) or high (affirmation) in value relative to other experiences.

In Q-Learning, the choice as to which state to enter as an action in an experience may be determined randomly, since the agent will learn and ultimately converge on an optimal policy without needing to refer to a representation of this optimal policy during its explorations of the domain. Hence, Q-Learning implementations may be labelled “off-policy” (Poole and Mackworth 2010, 470). Additionally, Q-Learning implementations do not require knowledge of what particular states represent in terms of the specific pathfinding problem they model. Any pathfinding problem that can be framed in the abstract terms of a domain of states, a set of possible actions, and a set of consequential reward values may be approached using Q-Learning, (Sutton, Barto, et al. 1998, 3.1) whether states represent a maze, a series of abstract choices, or any other broadly-defined set of spaces or conditions.

When the agent takes an action, it experiences a numerical reward which contributes to its knowledge of the domain. The values of particular actions are all stored in a table known as the Q-function (Here abbreviated “QF”). (Poole and Mackworth 2010, 472) See Table 5.1 for a depiction of two example QFs. The first is presented in its initial state prior to learning, and the other QF is configured in a state after learning has begun. The rows of the QF comprise the set of all possible states in the domain which can be occupied by the agent. States are denoted with an s symbol. The columns represent the total set of possible destinations to which the agent may transition through undertaking an action. An action is denoted by the symbol a . Often, when the first experience of an agent occurs, the QF cells are initialized to zero, or they are null. QF cells corresponding to a valid pairing of state and action—that is, one where the action leads to a state that is accessible from the original state—are set to zero. All other cells are null in value (here marked by hyphens), since such actions are not possible in the particular map of states in the given domain. A QF cell can be denoted based on the row and column it occupies. For example, an agent moving from state s_5 to state s_6 is referred to as $Q(5, 6)$.

In the Figure 5.1 example, there are six states with arrows showing permissible routes of travel. An agent in state s_1 can freely move to state s_2 , and hence the cell at row one, column two is initialized with a zero value. As learning occurs, this zero values may be modified to reflect the agent’s growing understanding of the value of

Q	a_1	a_2	a_3	a_4	a_5	a_6
s_1	-	0	-	0	-	-
s_2	0	-	0	-	0	-
s_3	-	0	-	-	-	0
s_4	0	-	-	-	0	-
s_5	-	0	-	0	-	0

An initialized Q-function.

Q	a_1	a_2	a_3	a_4	a_5	a_6
s_1	-	4	-	0	-	-
s_2	0	-	0	-	0	-
s_3	-	0	-	-	-	5
s_4	0	-	-	-	0	-
s_5	-	-2	-	0	-	5

A Q-function after some learning has occurred.

Table 5.1: Two depictions of the Q-function for the domain illustrated in Figure 5.1. The top table is initialized and displays no knowledge of the domain. The bottom table is populated with some nonzero values formed through the agent’s experience of differential rewards. Higher values correlate with higher experienced rewards. Rows correspond to states in the domain, columns refer to available actions from those states. Hyphens imply non-permissible state-to-state transitions.

taking action a_2 from state s_1 . Of course, the null hyphens on cells $Q(1, 3)$, $Q(1, 5)$ and $Q(1, 6)$ are present in reflection of the fact that the agent cannot reach any of these three states from state s_1 .

As has been noted, learning is a consequence of the agent encountering differential reward or punishment values based upon actions it takes during a series of episodes. Reward values are stored in table known as the “Reward Table” (Here abbreviated “RM”). Examples of such values can be seen in Table 5.2. Individual reward values are mapped into the RT in a manner corresponding to the QF, with the baseline reward for an agent transitioning from state s_1 to state s_2 located at call $R(1, 2)$. Higher values in the reward table are associated with more desirable actions. In the example, the agent is rewarded with a high positive value of 100^3 for moving from s_3 or s_5 to s_6 to reinforce its reaching the terminal state s_6 . Likewise, actions from s_3 and s_5 that are directed away from the terminal state s_6 are here punished with relatively low values of -50.

While reward values in the RT may remain fixed as the agent explores the domain, the QF evolves as the agent acts. QF values are modified by solving a variant of the “Bellman Equation,” which was first described by Richard Bellman in 1957 (Bellman 1957, 1966). The rule manifests in equation 5.1 (Sutton, Barto, et al. 1998, 6.5) that can be used to update individual QF values every time the agent undergoes an experience:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)]. \quad (5.1)$$

Symbols appearing in this equation are described as follows:

- $Q(s_t, a_t)$: The QF value for the pairing of state s_t and the action a_t , as it appears at time t . Every time the agent takes an action, t increments upwards by one (arbitrary) unit. Each individual value of the QF (that is, the value for a particular pairing of state and action) may be defined as the sum of reinforcements received when performing that action from the given state, and then following the policy thereafter. (Harmon and Harmon 1997, 10)

3. Rewards values are quantified relative to one another, but their absolute value is otherwise arbitrary.

R	a_1	a_2	a_3	a_4	a_5	a_6
s_1	-	0	-	0	-	-
s_2	0	-	0	-	0	-
s_3	-	-50	-	-	-	100
s_4	0	-	-	-	0	-
s_5	-	-50	-	-50	-	100

Table 5.2: A Reward table for the domain in Figure 5.1.

- r_{t+1} : The RT value corresponding to the QF cell that the agent currently occupies and the action that it chooses to take while there. It contains a numerical value proportional to the degree of reward associated with the agent's current state and action. It is the job of the designer of the reward system to populate these values in a way that reflects the goals of the agent. (Harmon and Harmon 1997, 3) That means that if a particular action is desirable, it will have a higher value than other actions that lead to less desirable or more neutral paths.
- $\max[Q(s_{t+1}, a)]$: The highest QF value among all possible actions that may be taken from the agent's chosen destination state. If, for example, the agent were to take an action moving it into state s_2 , and s_2 has three exits with current QF values of 3, 4, and 5, the value of this component of the equation would be 5.
- γ (*i.e.* "gamma"): The "discount factor" which weights the agent's ability to look ahead during pathfinding. (Even-Dar and Mansour 2003, 1; Poole and Mackworth 2010, 402) This value may be in the range $0 \leq \gamma < 1$, and reflects the degree to which possible reward values for actions taken subsequent to the current experience impact upon current learning. Discounted rewards mean that the effects of the current action on the QF may be augmented by the value of rewards from future actions, but the further forward the projection, the less its impact. (Watkins and Dayan 1992, 280) For values of γ greater than zero, convergence may not be assured, but learning occurs at a faster rate. (Glorennec 2000, 24) A γ value of zero causes the agent to ignore future rewards, slowing learning, but making convergence theoretically possible. (Poole and Mackworth 2010, 402)
- α (*i.e.* "alpha"): The "learning rate", which governs how quickly QF values are updated over a series of experiences. (Gasser 2009) Learning rate values occur in the range $0 \leq \alpha \leq 1$. Every time a QF value is overwritten per Equation 5.1, the α value acts to reduce the degree to which new information overwrites the current value of the relevant QF value. For higher alpha values, rewards (current and future) have a greater impact on QF values.

Every time that the agent engages in an action, this equation is solved and the corresponding QF value is updated, which is essentially the embodiment of the agent learning. Over many episodes, QF values will become ever more refined and the agent's knowledge more consistent with a hypothetical ideal configuration of the QF, and such convergence will occur so long as all state-action pairs in the QF continue to be updated through the actions of the agent. (Sutton, Barto, et al. 1998, 6.5) Theoretically, the agent's actions will refine the QF values to their ideal state over an infinite number of episodes, achieving convergence and revealing an optimal policy (Watkins and Dayan 1992). However, executing an infinite number of episodes is essentially impossible, so at some point an agent's updates of the QF must be halted. If enough episodes have been undertaken by the agent, the QF has been refined to the point that it can reveal the desired optimal policy, a situation which implies convergence is sufficient.

Revealing this optimal policy is the object of Q-learning (280). The optimal policy may be revealed after the agent has concluded its episodic traversals of the domain. To accomplish this, the agent is made to traverse the domain one final time. In this final episode, however, the agent does not use the policy it used during the learning phase, a policy that has no relation with the policy that is now being revealed (a reflection of the "off-policy" nature of Q-Learning described above). Instead, the agent, beginning in the origin state, traverses the domain by selecting in each state it encounters the action associated with the highest QF value. When this is done, and if the Reward Matrix has been populated with values that accurately reflect those actions' desirability, the agent will trace a path that ends in the destination state, maximizing the total QF values of all actions. This pathway constitutes the optimal policy, and reflects the quality of the agent's experiential learning. (281)

5.2.3 Generality of Domains

So far I have been referring to the space inhabited by the agent as a domain consisting of multiple interconnected states, but these words are general terms. One benefit of implementing pathfinding with Q-Learning is that it is agnostic with regard to what might constitute a state. The nomenclature is suitably general because Q-learning's applications can be broad—pathfinding with Q-Learning is technically

feasible no matter what a domain or state represents, so long as that network manifests a collection of states which through the agent may traverse and, by so doing, learn rewards that reveal the optimal policy. This is a property that Q-learning shares with other Reinforcement Learning paradigms. (Harmon and Harmon 1997, 3)

As such, these states can represent many hypothetical environments. One model might map states to rooms in a house. Each discrete room may be considered a state, and each state may connect with a subset of the other locations (such as through doorways or stairs), creating a series of permissible pathways from room to room, or equivalently from state to state. The domain is the collection of all possible locations or states that the agent might occupy, which, in our example, would be the house containing the rooms. In such a case, a typical objective for a Q-Learning implementation might be to determine an optimal pathway between two states—that is, from an initial room to a target room such as finding the quickest route from the front entryway to the main upstairs bedroom. The Q-learning system designer might reward the agent exiting the front entryway in the direction of the stairs to the upper floor, and punish the agent for taking a detour into the garage, which would increase its distance to the bedroom.

However, domains need not be mapped onto physical space. For example, Christopher Amato and Guy Shani evaluated a Q-learning-based approach for switching between high-level strategies in the video game *Civilization IV*. (Amato and Shani 2010) In their application of Q-learning, the eighty-one states represented a set of strategies that could be used to defeat an automated opponent, and an action represents a shift in strategy. (5) This example illustrates the fact that a pathway may not necessarily be rendered in physical space. Any series of states through which a pathway can be found may potentially be modelled in a Q-learning implementation. In the case of my own application, Q-Chords, that space maps musical harmonies.

5.3 Towards Q-Chords: Kristopher Reese’s Approach

In my program Q-Chords, states are equivalent to three-note chords and the domain comprises the set of all possible chords that maybe possibly occur in a transitional chord progression, including all possible combinations of three pitch chromas. This includes all triads of all scales, plus all possible non-triadic, three-note chords.

The Q-Learning algorithm is deployed to generate a sequence of chords which are drawn from the total set of chords that populate the domain, forming a harmonic progression (i.e. a sequence of three-note chords) linking user-input start and destination chords. This approach elaborates upon a Q-Learning framework described by Kristopher Reese in his master’s thesis 2011, and further developed in a 2012 article. (Reese 2011; Reese, Yampolskiy, and Elmaghraby 2012) My own system was presented at the joint ICMC/SMC conference in Greece in 2014, and published in the conference’s proceedings (Cullimore 2015). Because Q-Chords is partly a refinement of the approach of Reese et al, I will first discuss their approach to generating music with Q-Learning, and then contrast it with my own implementation.

Reese’s approach to generating tonal music applies Q-Learning to the problem of generating music “on-the-fly” during gameplay. (Reese, Yampolskiy, and Elmaghraby 2012, 131) Reese names his application “MuseGEN”, and positions it as an “engine that contains all of the primary functionality for generating music.” (Reese 2011, 46) According to Reese, MuseGEN contains three primary components: a database of possible scales, a rhythm generator, and a module he called the “Tonal Harmonic Progression” generator, which incorporated programmatic classes based on Q-Learning used to generate harmonies. (46-47) Since my own goal was to create generalized chord patterns that could be orchestrated into transitions via other applications, I built upon his third, harmony-centred module when developing my own approach to generating chord patterns.

To create a mapping of chords to a domain of states that would function in a Q-Learning application, Reese draws upon the work of Dmitri Tymoczko. Tymoczko proposes a system for organizing chords within a geometric space known as an “orbifold.” (Tymoczko 2006, 72). While a detailed discussion of the mathematical features of an orbifold is beyond the scope of this dissertation, Tymoczko’s application for chords of three or fewer notes is not hard to understand. Tymoczko proposes a geometrical representation of pitch space in which chords of n pitches may be mapped into an n -dimensional space, wherein each possible chord (comprised of a specific and unique combination of pitches) exclusively occupies a specific coordinate within the domain. (73) See Figure 5.2 for a reproduction of Tymoczko’s mapping of pitch space for two-note “chords”. While harmonies of fewer than three distinct pitches are not typically labelled “chords”, I use this terminology for the sake of simplicity and to




Figure 5.2: Tymoczko’s two-dimensional chord space (Tymoczko 2006, 73). Each point in this domain represents a unique two-note chord. Scanning left-to-right corresponds to transposition, scanning vertically increases or decreases the size of intervals between notes in the chords. The table wraps at the edges, with the upper-left corner aligning with the bottom-left.

more clearly link the concept of the example table to others featuring more complex chords. Each entry in the table consists of two numbers which represent the pitches in the chord, where any arbitrary pitch may be assigned the value zero, and other pitch classes are consequently labelled according to the number of semitones they are above the first pitch.

Note that this system of grouping and labelling notes derives from set theory, which has been described as “the foundation of all branches of modern mathematics” (Wierman 2010, 11). In set theory, objects are grouped into classes or collections called “sets” (11). Once these objects are collected into sets, the sets may be processed by functions. A function may be defined as “a mathematical abstraction of a consistent machine” (16). In other words, when a specific object is entered into a particular function, an identical project is output by that function (16). The generality of what can be put into a set, and thus what can be acted upon by a function, has allowed music theoreticians like Tymoczko to apply set theory to the representation and manipulation of music. For example, the system described above for representing chords as integer values mapping onto pitch classes is an application of set theory. Music theory is associated with many applications of set theory to tonal and atonal music, such as John Rahn (Rahn 1979) and Allen Forte (Forte 1964). I deploy a similar system of representing pitch sets in each of the subsequent works described in this thesis, as will be described in later chapters.

Every possible combination of two discrete pitches (including unisons) appears exactly once in this two-dimensional space. More importantly, the positioning of these two-note chords follows a regular pattern. For every chord in the space, the neighbour to the right is the same chord transposed up one semitone. If the pitches form a perfect fifth, for example, all pitches within that row will also form a perfect fifth, albeit ascending in root pitch by one semitone as the row is scanned rightward. Furthermore, when ascending in the vertical dimension, the interval between the two pitches increases. For example, a row of notes a semitone apart can be observed, underneath a row of notes a two semitones apart, which are in turn underneath a row of pitches forming an interval of three semitones. (Tymoczko 2006, 73) In other words, proximity of chords in this table is correlated with their similarity in tonal space, specifically with respect to their interval degrees and root pitches.

Tymoczko describes the configuration of this space as being akin to a Mobius

strip, in that the grid wraps around at the edges, allowing endless movement between intervals and degrees of transposition. (Tymoczko 2006, 73) Reese et al. observed that, were an agent to move around this space by transitioning from coordinate to coordinate, the series of chords associated with each step would constitute a chord pattern, a sequence of chords. (Reese, Yampolskiy, and Elmaghraby 2012, 132-133) See Figure 5.3 to observe how movement in Tymoczko's orbifold may be transcoded into chord sequences. Reese et al note that more complex chords of more pitch classes can be mapped into an orbifold simply by adding dimensions—a space representing three-note chords can be represented within a three dimensional topology, four-note chords in a four-dimensional topology, and onward for chords featuring larger numbers of pitches (133).

Reese sought a novel method for generating music that drew upon the order of these systems of pitch relationships. His approach was to use Q-learning as a framework for exploring the domain of harmonic tonality and generating musical events. As such, the domain became the chord space as in Figure 5.3 and its constituent states became the chords contained within. The agent's travels reveal a pattern of chords that correspond to the states that the agent passes through on its travels. However, Reese notes that while the Q-learning framework provides a means to map an agent's travels through the domain onto a sequence of chords, it does not inherently specify a policy to guide the agent's explorations through the domain—that is, to answer the question as to where the agent should go next as its action. (Reese 2011, 66) Without such a policy, Reese could not create a system for navigating the chord space based upon Q-Learning alone.

Reese's solution was to use a velocity vector to represent the travels of the agent, and create a behavioural model governing the movement of that vector. (67) However, at this point Reese's framework diverges from my own application. While Reese chose to replace actions with vectors with specific magnitudes and directions, and create a policy governing the specification of those vectors as a means of directing the agent, I chose to manipulate the structure of the chord space itself. Reese used three-note chords in his explanation of his approach, partly because such three-dimensional orbifolds are easier to visualize than four- or five-dimensional forms that may be associated with harmonically-rich music such as some jazz (Reese, Yampolskiy, and Elmaghraby 2012, 133). Regardless, a three-dimensional orbifold does provide ample




Figure 5.3: Reese's representation (Reese, Yampolskiy, and Elmaghraby 2012, 132-3) of a chord pattern as mapped in a two-dimensional orbifold (above) and in musical notation (below).

opportunity to develop a policy for exploration of chord space.

5.4 Q-Chords: A Novel, Q-Learning-based Chord Transition Generator

Unlike Reese’s MUSGEN, my own system Q-Chords, does not use vectors, instead deploying other methods to create musical variation, described below. My goal is also quite different. I was not concerned with generating potentially-endless compositions, but instead had a very strictly-defined goal: to create short chord sequences that could effectively serve as bridges between musical material in different keys⁴. This construction is described by Tom Sutcliffe as a “dynamic harmony,” a form of modulation that moves from one tonal centre to another through a sequence of chords (Sutcliffe 2018, Section 7.1), and the chords that may be included in such a progression are derived from the triads formed by the root notes of the scale (Section 2.2), and so are diatonic in nature. A harmonic progression, he says, can use pivot chords, a technique introduced earlier to which I will return in the description of the functioning of Q-Chords.

To this end I propose three main alterations to Reese et al.’s model for chord sequence generation:

- To increase the efficiency of the Q-Learning process by reducing the number of chord states in a domain.
- To generate chord sequences that avoided what I subjectively assessed as unnecessary dissonance, by selectively excluding chord states that diverge from the tonal space defined by the start chord or goal chord within a domain.
- To initiate generation of transitional chord sequences early by concluding the exploration phase and revealing the optimal policy before prior its approaching

4. The choice to limit Q-Chords to construction of transitional chord sequences, and not other, longer forms of progression, is not due to a limitation of Q-Learning. Any type of progression—chordal, melodic, timbral, or otherwise—may be operated upon through a Q-Learning application. I restricted my project to a narrow goal because the project was originally designed for a course project that necessarily involved limiting its scope. Conceivably, because the number of domains in a Q-Learning application need not be limited, much larger domains are possible providing a concurrent increase in the possible complexity of the web of interconnections between states. As such, Q-Chords points towards the feasibility of more complex musical applications of Q-Learning, rather than serving as an indicator of the inherent limits of such an approach.

of convergence. This effectively creates greater variation in generated chord sequences while retaining their tendency to conform with desired tonal qualities.

Q-Chords is not intended to generate fully-realized compositions. Rather, its primary aim is to serve as a compositional auxiliary, a tool for automatically creating bridging chord sequences between musical cues in major tonalities. This system could be extended to work with minor or modal tonalities, however, that was beyond the scope of this project as I deemed a successful proof-of-concept could be constructed using only major keys. The program's interface grants the user an ability to control the expressive and tonal qualities of generated music. In my experience, when a composer creates music in a certain key, they may be aware that some pitches and intervals are more common and useful than others, particularly when attempting to control aesthetic qualities such as emotional valence and intensity. The tritone, for example, is an interval of six semitones distance, and it typically and effectively produces dissonance; this aspect of the interval makes it more useful to composers aiming for dissonant effects than consonant ones. Similarly, in Q-Chords, the RT acted as a biased filter which could eliminate chords that did not conform to my aims—for example, if I wished to suggest a major tonality, I could use the RT to promote the incorporation of diatonic triads (in relation to the start or end key of the sequence) within generated progressions by rewarding actions that lead to such a diatonic triad. Despite biasing the RT with values that strongly promote major (and to a lesser degree, minor) chords, Reese et al.'s generated chord sequences still can exhibit a high degree of chromaticism, since all possible chords are represented as states in their framework. A composer writing in a Western Tonal idiom such as the classical symphony or conventional pop song may know that there are certain chords that are much less commonly observed in their particular style of music, and that in choosing harmonies for their music, such chords might be safely ignored. Thus, in Q-Chords, I streamline the chordal domain by eliminating chord choices that strongly depart from the tonality defined by the start and goal chords. This biasing of the RT increases the chance that progressions will include diatonic chords—that is, triads built from roots originating in the major scale associated with the start or end key—and that pivot chords, defined as diatonic pitches found in both the source key and the destination key (Goldenberg 2018), are even more likely to appear.

Encouraging the appearance of diatonic triads through biasing the RT was intended to promote the appearance of familiar harmonic progressions from Western Tonal music. Tom Sutcliffe explains that the function of progressions using dynamic harmony is to “provide a sense of moving forward” (Sutcliffe 2018, 2.1), a goal compatible with my own objective of automating the composition of cue-to-cue transitions for video games. Sutcliffe points out that the use of chord inversions (where the lowest pitch of a chord is not its root) can weaken this sense of forward momentum (2.1). This is concern, as I chose to keep the generated chord sequences in an octave range, which arbitrarily inverted some chords, but not others, during playback. Other researchers, including Reese et al., use algorithms to permit voice-leading, a possible future direction for my own development of Q-Chords.

5.4.1 The Domain in Q-Chords

To improve the efficiency of its pathfinding, and to allow the user input into the expressive qualities of generated music, Q-Chords constructs a domain out of a subset of all permutations of three pitches. The question of inversions arises here, since a chord of three pitches may be assembled in various configurations depending on which octave each of the three pitches occupies. Different octave placements influence the ascending sequence of pitches in a chord, and a chord formed of the same three pitches may not be functionally identical to another chord featuring the same three pitches in a different ascending order. My approach to representing chords draws upon Reese’s framework, which treats chord pitches as pitch chromas rather than pitches with a specific fundamental frequency. Pitch chromas represent positions in the twelve-note chromatic scale, and unify all notes in octave relationships. They are a means of organizing pitch complementary to pitch height, in which pitches are defined in proportion to their fundamental frequency. (Bachem 1950, 1873) All notes labelled C, for example, are said to belong to the C-chroma, regardless of their octave or differences in fundamental frequency. The same unity is observed for all notes of the chromatic scale, so in WAM there are typically twelve commonly-recognized chromas.

As in other applications of set theory, utilizing chord designations based upon pitch chroma allows chord representations to become generalizable. Reese gives the

example of an A-major chord, $A - C\sharp - E$, which, when converted to chromas with the A-chroma as root, can be represented as the chord $[0, 4, 7]$. (Reese 2011, 56) Reese notes that in MUSEGEN, all permutations of a chord of specific chromas are functionally equivalent, so that the chords of, for example, $[4, 0, 7]$ and $[7, 0, 4]$, are indistinguishable in the Q-learning domain. (56) Since chord inversions of this type are actually functionally distinct in music theory, projecting differing expressive qualities, Reese implemented a voice-leading algorithm that could convert chord sequences formed of pitch chroma to specific pitches and chord inversions so to ensure generated chord sequences might flow smoothly (69). The term “Voice leading” refers to “rules about melodic motions of voices involved in harmonic progressions” (Hutchinson 2017, 26.1). This is necessary because harmonic progressions involve sequences of chords that are formed of multiple pitches each, and these pitches can almost always be rearranged as inversions that will take on different characters when used in alternate constructions of a particular harmonic progression. Voice leading governs the melodic lines that form as individual voices (such as the soprano, alto, tenor, and bass parts of some choral works) sing their individual musical lines. The practice of voice leading is primarily concerned with maintaining the independence of the parts, and promote melodic motion that is singable or otherwise appropriate for an instrument who contributes that line to the progression of chords (26.1). The result was that chord inversions were rendered redundant in their form in the domain, and could be represented by a single tuple pitch chromas, considerably reducing both the size of the domain and the efficiency of pathfinding trials (Reese 2011, 58).

My own approach to reducing domain complexity was based upon the application of my understanding of music theory and my personal opinion regarding the utility of particular chords, based upon my experience as an autodidactic composer. That is not to say that I am in any sense unaware of core music theoretical concepts; I have a knowledge assembled from my own informal study of music theory, and a long and varied experience in composing and listening deeply to music. To this end, rather than devising a system for converting chroma-based chords into functional chord patterns, I chose to restrict the Q-learning domain to the set of chords that I subjectively felt were most consistent with the overall aim of Q-Chords, namely to generate chord transitions promoting consonance and preserving major or minor tonality. I saw these transitional patterns as a skeletal structure—raw harmonies which could, with further

programmatic development, be embellished with melodic and rhythmic figures. As such, voice-leading was not required. Instead, I restricted the domain of included chord-states, first assuming that the start and goal chords would be major chords (which I believe are common in tonal film and video game scores, although this is not to exclude other frequently observed chords such as minor or diminished varieties). Since major chords comprise a subset of the notes within the five-note (pentatonic) scale (for example, the C-major chord, with notes $C - E - G$, represents a subset of the C-major pentatonic scale, $C - D - E - G - A$), it is possible to restrict the domain of chord states in which Q-Learning will occur to include only chords that may be built from the pentatonic scales of the start chord, the goal chord, or a combination of the two. An illustration of such a domain (based on C-major) is shown in Figure 5.4, which shows a complete domain for chords of $n = 2$ (i.e. two-note intervals) with excluded chords greyed out.

Since this domain is significantly smaller than the full, chromatic domain, it can be expected that the agent will, during an episode, tend to take fewer steps to reach the goal state. The agent’s pathway is randomly directed and the number of choices of action that the agent may have at any given state is reduced, meaning that each step through the domain is more likely to reach the goal state.

A domain based upon the pentatonic scales of the start chord and goal chord is constructed as in Figure 5.4. Any chord states within this domain consist entirely of pitches that are found in the pentatonic scales based on the start and goal chords. Any chord states that contain pitches not found in either of the two pentatonic scales are excluded. If the start and goal chords are the same, then all chords in the domain will be built exclusively from the five pitches in the pentatonic scale based on that chord. If the pentatonic scales do not overlap at all, then the domain will feature chord states that draw from a set of ten pitches. Figure 5.5 illustrates how this process accepts pitch chromas for a domain generated from the C -major and $E\flat$ -major tonalities.

Note that such blending of chords originating from two relatively-distant keys, a state of bitonality may be achieved—think of Stravinsky’s notable “Petrushka” chord, used in his 1910 ballet of the same name, which layered the triads of C major and $F\sharp$ major to greatly dissonant effect⁵ . Such dissonance has a source in the distant separation of the two component keys, and bitonality has historically been

5. Described on the Take7 blog: <https://www.take7music.com/bitonality-in-classical-music/>

C-C		C#-C#		D-D		D#-D#		E-E		F-F		F#-F#
	C-C#		C#-D		D-D#		D#-E		E-F		F-F#	
B-C#		C-D		C#-D#		D-E		D#-F		E-F#		F-G
	B-D		C-D#		C#-E		D-F		D#-F#		E-G	
A#-D		B-D#		C-E		C#-F		D-F#		D#-G		E-G#
	A#-D#		B-E		C-F		C#-F#		D-G		D#-G#	
A-D#		A#-E		B-F		C-F#		C#-G		D-G#		D#-A
	A-E		A#-F		B-F#		C-G		C#-G#		D-A	
G#-E		A-F		A#-F#		B-G		C-G#		C#-A		D-A#
	G#-F		A-F#		A#-G		B-G#		C-A		C#-A#	
G-F		G#-F#		A-G		A#-G#		B-A		C-A#		C#-B
	G-F#		G#-G		A-G#		A#-A		B-A#		C-B	
F#-F#		G-G		G#-		A-A		A#-A#		B-B		C-C

Figure 5.4: An example of a pentatonic chord space built on the pentatonic scale with root *C*. Chords that contain pitches not included within this pentatonic scale are greyed out, and would not be included among the states of a domain built upon a C-major chord. Reproduced from (Cullimore, Hamilton, and Gerhard 2014, 337)

associated with controversy, with some arguing that the concept is impossible in tonal music, since by definition tonal music has a single pitch root, never two (Hamamoto, Botelho, and Munger 2010, 2010). Q-Chords, by forming harmonic progressions that blend pitch sets drawn from the diatonic scales of any two roots, naturally strays into the realm of bitonality in certain configurations. This is not an error, but a design feature intended to permit the system to help video game composers write as freely as possible, in any key, secure in their knowledge that Q-Chords could provide a harmonic progression linking it with any other, regardless of how distantly the two chord roots are situated in tonal pitch space.

I use pentatonic scales built on the chords roots of start and destination chords because the pitches of pentatonic scales can form a large number of consonant intervals when played simultaneously. There are several perfect fourths, perfect fifths and major thirds. Pentatonic scale pitches also exclude certain dissonant intervals, such as tritones and minor seconds. Thus, music constructed out of the notes of a pentatonic scale tends to take on a pleasingly consonant aspect. This is a subjective choice that reflects my own taste and biases towards composing consonant music, and my appreciation of other film and video game composers' work that is similarly formed. Other developers might choose other means of culling the domain, for example if they desired dissonant sequences instead of consonant ones. As such, this aesthetic choice I made arguably led to a key event in the development of my doctoral creative work: an element of my compositional ideals embedded in the form of program code. I had begun to enter the machine as a mirror of my agency, because the machine, in executing the code I designed, exhibited a trace of my composerly values. The culling system could have prioritized dissonance, or variety, or only chords containing the pitch $C\sharp$. That it functioned to promote consonance in its chords means that the output of Q-Chords is biased in the same way my notated compositions are—one can listen to one of my *We*, the *Artificial* tracks and recognize the same value in the output of Q-Chords. If my overarching goal as a doctoral student was to meet other human agents in producing an algorithmic composition, this goal was not met here. But its precondition, to find myself in programmatically-produced music, was achieved. In this manner, Q-Chords can be identified as the beginning of my exploration of agency in music, in that it was the first instance where I found myself in my algorithmic music.

All Pitches	C	C \sharp	D	E \flat	E	F	F \sharp	G	A \flat	A	B \flat	B
A-minor	C		D		E			G		A		
E \flat -Major	C			E \flat		F		G			B \flat	
Exclude		C \sharp					F \sharp		A \flat			B

Figure 5.5: Exclusion of pitches based upon the start and goal chords. Only chords built entirely from pitches that are represented within the pentatonic scales built from A minor and E \flat -major are included in the domain. Reproduced from (Cullimore, Hamilton, and Gerhard 2014, 337)

The reward system, too, can have a strong influence on output musical structure. With a thoughtfully-designed RT, a generated chord sequence can embody a variety of different stylistic and aesthetic aims. Reese et al. construct their RT as a set of states, with a set of allowable actions for each of those states. (Reese, Yampolskiy, and Elmaghraby 2012, 133). RT values in their system were designed partly by trial-and-error-based experimentation, as well as some “basic assumptions” about tonal music. (135) They settled on rewarding major and minor chord-states highly, as well as a bonus reward for chord states that contained pitches conforming to the scale of the composition being constructed. The Q-Chords system uses a different approach to setting reward values, with a dynamically-populated RT generated based upon user input.

Table 5.3 illustrates how RT values may be generated by Q-Chords, depicting a subset of a hypothetical RT built to support the generation of a transition from C -major to $E\flat$ -major chord states. The construction of the RT begins by analyzing the subset of chord states that have been admitted into the domain as shown in Table 5.5. It examines each chord state and its associated actions, and notes the pitches the agent will encounter while undertaking that transition (i.e. it considers the pitches of the chord associated with the agents’ current state and those of the state into which it is transitioning as its action). It then analyzes this set of pitches to determine whether it conforms to a pentatonic scale built from either the start or end chord of the chord pattern being generated. For example, in the RT in Table 5.2, there are two pentatonic scales upon which the domain is built, with roots $E\flat$ and C . If the combined pitches of a particular state and action do not all belong to one of the two pentatonic scales, there is a potential for dissonant intervals to occur and thus such state-actions pairs are rewarded with a punishment of -50 . An example of such a pair would be motion from state $E\flat - F$ to state $C-D$ (note that this set of notes contains the potentially-dissonant interval $D - E\flat$). On the other hand, potentially-consonant motion within one of the two pentatonic scales (e.g. $E\flat - F$ to $C - E\flat$) is rewarded with 100.

Note that between the pitch sets of the two featured pentatonic scales there exists a pair of pitches that appears in both, that being $C - G$. Because this interval is found in both scales, it might act as a “pivot” point between them, serving as a combo ground sustaining smooth transitions between the C -major tonality and the

E^b -major tonality. A state/action pair that results in the agent reaching this pivot state is given an additional, generous reward of 100.

The goal chord may also be heavily negatively biased (with a very low reward value), to ensure that motion in chord space avoids this chord. I made this aesthetic choice in my code design since I recognized that the goal chord would always be added to the end of the generated transitional sequence after it was generated, and biasing against its appearance in the transitional sequence could help the transition avoid sounding repetitive. To create the sense of a cadence leading from a generated chord transition to the goal chord, the dominant chord in the key of the goal chord is inserted immediately before it in the sequence. Because this dominant chord (that is, the chord with a root situated on the fifth scale degree of the key) is highly compatible with the tonal space inhabited by the goal chord, its insertion into a sequence as the penultimate chord I feel creates a heightened sense that the sequence is building towards the goal chord in a manner consistent with norms of tonal music.

5.4.2 Hurrying the Agent

If enough episodes are undertaken, a QF will approach convergence, through which an optimal path may be identified linking the start state and goal state. In dynamic music, however, it may be advantageous to have some uncertainty in the structure of a chord transition. The video game player might hear a transition between two themes dozens of times, depending on the game state, player behaviour, or sound design, and so randomness in the structure of a transition can help the music to remain fresh and unpredictable. At the same time, it may also be desirable for generated chord transition sequences to consistently conform to the tonal norms I have already described, an aim which may limit the scope of variability.

Reese et al. promoted variability in MuseGEN through a feature they built into the pathfinding algorithm occurring after the learning phase. When their agent made state-to-state transitions, it would identify the most beneficial action leading to a state based on the available QF values and move towards it. However, Reese included a stochastic inaccuracy, where the agent might overshoot or undershoot the target chord state. As explained earlier, actions in Reese’s model are substituted with velocity vectors, incorporating speed and direction values. During pathfinding, Reeses’ vectors

R	A1 C-D	A2 C-E \flat	A3 C-G	A4 E \flat -F	A5 A-B \flat
S1 C-D		-50	+200	-50	-50
S2 C-E \flat	-50		+200	+100	-50
S2 C-G	+100	+100		+100	-50
S4 E \flat -F	-50	+100	+200		-50
S5 A-B \flat	-50	-50	+50	-50	

Table 5.3: Example of reward values generated by Q-Chords based on user input. Units are arbitrary. Reproduced from (Cullimore, Hamilton, and Gerhard 2014, 336)

may have their speed value increased or decreased, such that the vector might fall one or two pitches short of the target pitch, or overshoots it by one of those amounts. (Reese 2011, 67) The result is that the path traced by the agent as it begins to follow its optimal policy will not follow the path of greatest reward, but rather will tend to wander into less desirable, but possibly more musically-interesting, chord states.

In Q-Chords, my means of promoting variety occurs before the optimal policy is converged upon: the algorithm is allowed to run for only a relatively small number of episodes, such that the agent deliberately fails to achieve convergence. Since QF values are updated as a consequence of the random actions of an agent within the chordal domain, lower numbers of iterations will result in QF values that are more heavily biased in areas that happened to be more frequently travelled by the agent. The less the agent has travelled through the domain, the more likely it is that the agent's random traversals of the domain will develop some areas of the QF over others. Thus, I added an input in the Q-Chords interface for users to control the number of iterations that the agent undertakes before pathfinding generates the transitional chord pattern.

This is not to say that cutting off the agent's explorations before convergence will result in chaos. Q-Chords' generated sequences largely avoid dissonance, because the pruning of states occurring prior to learning serves to restrict the degree of dissonance possible. For this reason, tonal consonance can often be achieved without the requirement of extensive exploration of the domain. I feel this to be an advantage of my system over MuseGEN, which Reese notes exhibits a tendency towards chromaticism (Reese, Yampolskiy, and Elmaghraby 2012, 135), and which, he admits, might benefit from further development in order to become more "pleasing to a typical listener" (136).

5.4.3 Application of Q-Chords within a Video Game

The original impetus for developing Q-Chords was to build a tool that could function within a video game scoring environment as an automated system generating two-measure transitions between (potentially-human-composed) cues. To apply Q-Chords to this task, it would require the composer of a video game score to provide a chord chart of each cue in their video game. This chord chart would, for every

measure, indicate the predominant chord underlying the music. Typically, this chord would be built upon the root of the composition's key. However, music of more than a little complexity might range around the harmonic circle, and even a piece scored in *C*-major might find itself centring on the strongly-related chord of *F*-major, or even, in particularly daring instances, travelling to the distant territory of *F* \sharp -major. Once that was provided to Q-Chords, the game might implement Q-Chords as follows:

Assume the player is exploring a serene environment accompanied by a pastoral musical theme *P*, written in *C*-major. At some point while *P* is playing, the game spawns a dragon, necessitating a transition to an energetic monster theme *M*, written in *E* b -major. Since the Q-Learning system described here is capable of transitioning from any major chord to any other major chord, it may be possible for transitions to occur starting from any point in *P*, so long as the chord that is currently harmonically predominant in *P*, at the time at which the transition needs to be instantiated, is known. There may even be no need to automate the map of underlying chord changes in theme *P*, where the composer chooses to enter this information manually. Such a generative transition system can support transitions from any point in a theme for which the underlying chord is known, since it can determine the initial chord in the Q-learning domain by choosing one that is identical to, or related to, the chord underlying the currently-playing music cue at the point at which the transition is instantiated. A further advantage would be that the composer need not restrict him- or herself to a narrow set of keys in a cue; since they are not restricted to preset layers or transitions that only function in one key, the composer might range widely through tonal space.

Of course, a chord pattern is not an arrangement, and the chord sequences that Q-Chords produces are best thought of as a skeleton upon which to drape the trappings of a fuller work. A harmonic map can be used to alter an existing arrangement, as will be demonstrated in the next chapter, and this is a primary reason why I consider Q-Chords a creative tool and not an artwork in its own right.

5.4.4 Contributions of Q-Chords

Q-Chords has several compelling qualities:

- Theoretically, transitions may be generated from any chord to any other chord,

meaning that composers are free to range as far as they want in tonal space.

- There is an inherent randomness in the generated chords, even though those chords can still retain an air of tonality, which helps them to remain effective while reducing the tendency of repeated transitions becoming stale and predictable.
- The reduction in the size of the domain (through excluding undesirable chord states entirely) serves to reduce the computational load on the system, increasing its efficiency and, potentially, allowing it to be implemented in near-real-time during execution of the game code.
- Biasing the reward table affects the tonal qualities of generated transitions, leading to the possibility that the reward table can be an expressive tool. It is possible, for example, to encourage transitions that focus on minor chords by simply increasing the reward values for them in the reward matrix.
- By building the reward matrix on two dimensions and considering the relationship between pitches in each state and each corresponding action, it is possible to base the reward for an action not just on the fit of the destination chord with the user's aims but also on the qualities of the transition between the current chord and the destination chord.

Through implementation of this system, it is theoretically possible for composers of adaptive music (such as that of video games) to focus their efforts on composing cues, rather than concerning themselves with composing a potentially large number of required transitions. They might feel freed to write more challenging or colourful music since they would not be obliged to confining their cues to a narrow range of keys. With creative manipulation of the reward table and choice of the number of episodes to iterate before generating a chord sequence, the composer can create desirable variability in their generated chord transitions while still maintaining their conformance with the tonalities surrounding the start and goal chords. As such, this application of the Q-Learning framework may serve both as a creative tool and a means of supporting the adaptive music composer's option of writing music of higher harmonic complexity.

5.5 Reflection on Q-Chords

In his article “Musicking Beyond Algorithms” (Bhagwati 2015) Sandeep Bhagwati explores how machines, long known to be able to amplify and supersede human physical capacities, are rapidly increasing in their capacity to amplify human intelligence, particularly in musicking contexts. While amplifications of human physical capacities, such as a hammer magnifying the arm’s strength, are often taken for granted today, Bhagwati notes that the idea of amplification of intelligence has been regarded with far more concern, or even outright fear. (373-374) Douglas Hofstadter conveys this sentiment in his analysis of David Cope’s *EMI* composing program. As he played upon the piano an emulation of a Chopin mazurka constructed by *EMI*, he reports experiencing a pang of emotion that he could not attribute to Chopin. Furthermore, he thought, were he unaware of the composition’s origins in Cope’s algorithms, it might have seemed to him to have been the product of a well-schooled human composer. (Hofstadter 2001, 4) Of course *EMI*, Hofstadter continues, is no human, and possesses no history granting it a well of experience upon which to draw for expressive effect. Indeed, *EMI* had no emotions whatsoever. Furthermore, Hofstadter noted that *EMI* did not simply construct incoherent, Frankenstein-like collages of a particular composer’s previous compositions. Rather, he observes that the program’s compositions seemed to project the style of the composer while disguising the origins of particular elements (42). Because *EMI* could so profoundly project the air of a human source in its most effective scores, Hofstadter’s early view on *EMI* and its music was the disconcerting sense that the program was robbing some of “the profundity of the human mind’s sublimity”. (6) Simply put, the machine’s output was too human.

At this early stage of my doctoral program, I largely shared Hofstadter’s concerns regarding relocating the site of compositional expression from human minds to machines driven by code. Even today I agree with contemporary statements that people like Tesla and SpaceX founder Elon Musk has issued—especially his alarming 2017 comment that Artificial Intelligence research may pose “vastly more risk” to human civilization than North Korea’s burgeoning nuclear capabilities.(Gibbs 2017) And it might be true, if computers were to achieve intellectual powers beyond our own and an autonomous independence that pitted us against them. These warnings have been

raised for decades. For example, the “singularity”, a hypothetical point where machine intelligence supersedes humans’, was starkly described by Vernor Vinge as the end of the “human era”. (Vinge 1993) Such fears of conflict with machines have entered the cultural vernacular as a compelling means to engage audiences—consider future wars depicted in the *Terminator* movies, pitting human against machine, or more recently, the unsettlingly human-like androids in the *Alien* franchise. In all of these cases, machines are villains whose human facade hides a potential for dark malevolence.

Yet beyond the machine-monsters of pop culture that warn us against blindly grafting human society to an ecosystem of machines, there are substantive alarms about computer automation being raised in academic and artistic communities as well. Writing in 2009, Margaret Boden observed that computers are already capable of what she terms “H-creativity,” the form of invention which created something novel that never existed before in history (Boden 2009, 24). She points to an AI program that generated a novel design for a chip that was subsequently awarded a patent, (24) and the classic drawing program AARON, which began as an attempt to grant a computer the capacity to create sets of marks that could be understood as an “image” by human viewers, but which eventually produced digital paintings exhibited in the galleries of the San Francisco Museum of Modern Art (1979) (Cohen 1995, 2-3).

Of course, I had been stunned by some of *EMI*’s original works in the style of Bach and Mozart. In this context, I felt that my contributions to computer music might come at the expense of my opportunities to express myself through specifying elemental musical structure in my traditional linear scores. I feared that by programmatically replicating my personality, values, and characteristic techniques for addressing musical challenges, I might find that I’ve lost control over my expression, and to own those expressions. There may be some justification to this fear. In June, 2019, for example, *Vice* magazine reported that the A.I. known as *AIVA* completed an unfinished piano piece by the composer Dvorak, 115 years after his death. The work was performed under the title *From the Future World* by the Prague Philharmonic in November, 2019 (Rudra 2019). *AIVA* is being marketed as a aid to composers. Its website claims that with *AIVA*, you may “create compelling themes for your projects faster than ever before, by leveraging the power of AI-generated music”⁶

. AI applications like this may encroach on compositional skills, and even render them redundant. I have argued this point in my article “The Virtuoso Composer and the Formidable Machine: A Path to Preserving Human Compositional Expression” (Cullimore 2015) in which I identify the pencil and paper as being the most difficult compositional tools to master, since they provide the fewest accommodations that counteract lack of skill. While a fuller discussion of this point will be delayed until the final chapter, it is relevant to the current discussion that I had always felt an unease with “offloading” compositional tasks to computers. I had taken to heart the attitude I had observed in my interactions with some WAM composers that eschewed composing in a digital environment. Even as I composed in Cubase, and was embarking on a doctoral program in which I would use code creatively to generate music, I felt a palpable sense that, by delegating portions of my creative action to an automated system, I was removing myself from my scores, and bestowing that space upon a machine whose capacities were impressive, but whose expressions were empty. The experience of Garry Kasparov, the chess grandmaster whose place in history was secured by his loss to IBM’s Deep Blue program, felt poignant to me. (Campbell, Hoane Jr, and Hsu 2002) Speaking during the run-up to his loss, Kasparov issued a dire warning:

To some extent, this match is a defence of the whole human race. Computers play such a huge role in society. They are everywhere. But there is a frontier that they must not cross. They must not cross into the area of human creativity. It would threaten the existence of human control in such areas as art, literature, and music. (Hofstadter 2001, 5)

Kasparov was defeated by Deep Blue in 1997 after the software had been revised in response to the software’s loss in a tournament during the previous year. (Campbell, Hoane Jr, and Hsu 2002, 59)

I was well aware of Kasparov’s defeat while composing Q-Chords, and wondered if I was on a path towards either coding something, or encounter another’s computer system, that might render my musical scores superfluous. But what would such a negation look like? At the time I felt that it would look like a draining of control, in the form of a tool that would amplify my musical intelligence to the point where

6. AIVA’s official website can be found at <https://aiva.ai>

my identity was interchangeable with other composers, due to a loss of my ability to control the expression of my ideas and emotions. Q-Chords, as the first interactive computer-based musical collaborator I had designed, reflected in its design my trepidation over letting the computer mimic or augment my own cognition-based compositional talents and behaviours. Consider the following features of Q-Chords:

- Interaction between user and software is segregated, in that users may affect variable values that impact upon generated transition content, but not after the process of chord sequence generation has begun.
- Harmonic patterns are treated as an output only—there is no ability to feed them back into processing once produced, and as such their appearance marks an end to the function of the program.
- Several significant musical choices are fixed into the functioning of the program and cannot be changed, such as the number of notes in a transitional chord, the length of the transition, and, more fundamentally, the Eurocentric harmonic language that forces all transitions to conform to Eurocentric norms. For example, a user cannot specify alternative tunings, nor the requirement that generated chord sequences begin and end in a major tonality.

These limitations reflect my desire to automate a compositional process that I would otherwise have to compose manually. As a consequence, Q-Chords behaves in a similar manner to the results that I would anticipate were I to compose a transitional chord sequence linking two major chords in my DAW software. Of course, the method is different—it is difficult to envision Q-Learning as mapping upon a cognitive process that occurs while I compose. Yet the issue here is not of the technique, but rather of the relationship between intent and product. When I compose, regardless of the particular cognitive processes that unfold, or the workflow of the software I use, my compositions tend to share certain features. I tend to value consonance, and consider my scores' accessibility to wider audiences one of its strengths. Q-Chords embodies such values—and insofar as it does so, it is a characteristic creative product of mine, no different from a linear score such as my 2016 score for *The Cabinet of Doctor Caligari*, or my drawings, which I feel project my desire to be clear and balanced in my expression, even when the subject of the drawing is abject in nature.

The key limitations of Q-Chords relate to this point, in that its output may most definitely be identified with me—my personality and musical values. Computers may bend not just to those who create the software, but also those who use it, and in fact any motivic force in the environment can be input into the running code via sensors, network connections, or internal structures that evolve within the code itself. The features I list above may actually be framed as limitations, since my tendency at the time was to lock down the functioning of the music generation engine, separating its function from users, and in so doing retain the authority over the output music that I felt might otherwise slip away. Like Eno’s *Bloom*, Q-Chords does not bend to alternate perspectives simply because it is designed not to. Each of the three design choices listed above represent my attempt to guard the code against users who wish to create in Q-Chords something that I wouldn’t expect or condone.

If this thesis is a study of composers’ agency and its transformation, then Q-Chords reflects my persistent desire to maintain control of music, to have agency over the product, and to remain elevated above others who are part of the process that unfolds when the program is engaged⁷. My agency remains privileged in Q-Chords, and neither the user, nor the code itself, can evolve a chord pattern beyond what I intended—consonant sequences of eight chords that begin on a major chords, end on a major chord, and, in-between these points, generate music that effectively functions in a manner that suits its simple transitional task. Without admitting other agencies, one is left with a rigidly-structured output that, despite its incorporation of randomness, is incapable of surprising me. It is thus similar to one of my notated

7. One might argue that Q-Chords becomes like a digital audio workstation such as Cubase—a software environment setting up conditions for users to create some musical object. I maintain that my agency is elevated in Q-Chords in a manner that may not be seen in Cubase or similar software. In Cubase it appears that its coders wish to provide flexibility to its users. For example, the interface allows a user to compose a symphony or a electronic track, with many points-of-entry where the environment might be modified through integrating third-party plugins or interface tweaks that personalize the composer’s experience. By contrast, Q-Chords allows minimal input on the part of the user, restricting them to the specification of variables that are input into the algorithm before it calculates the chord sequence. The users of Q-Chords have little agency over the output chord sequences, much less that a user of Cubase might. The result is that the process of using Q-Chords is less one allowing for the user’s self-expression and more one that constrains them to the role of a spectator, setting initial conditions and then observing the result, without means to intervene in the process in a manner that deeply reflects their self-expressive choices. As such, the user of Q-Chords seems to me like an audience member more than a creative partner, and for this reason I consider the function of Q-Chords, and the effect of the musical objects it produces, to remain largely under the control of my own agency.

or recorded scores: once the program is written, its potentials are fixed. By this I mean that it can have an effect on emotions or intellect when encountered, but it cannot extend its action beyond the possibilities of the code itself, which I would argue requires an interplay of agencies that redirect the action of the app.

As I moved forward from Q-Chords, I recognized that the program was very much a reflection of my fear of losing my primary expressive outlet (i.e. composing music) as a consequence of what I term a mingling of agencies that occurs within and around the execution of an environmentally-responsive computer. Yet to have authority over a programmatic unfolding may be considered a limitation of possibilities. In his 1958 doctoral thesis, Gilbert Simondon argues that the technical and the cultural, and between machine and human, is without foundation. (Simondon 2011, 1) Simondon draws on the idea of the “technical object”, which can be identified with the functional machine or tool, utilitarian in purpose and “anonymous and abstract” in its identity. (Dufrenne 1964, 115) Culture, Simondon states, treats a “technical object” as either utilitarian “assemblies of material” on the one hand, or as a robot, which may “harbour intentions hostile to man, or that they represent for man a constant threat of aggression or insurrection”. (Simondon 2011, 3) In both views, the human exists apart from machines, above them in the first case, and in opposition to them in the second. Simondon contributes a third path, wherein the human is actually “among the machines that work with him” (4). In the case of Q-Chords I was engaging in the first form of human-machine relationship, for want of warding off the real or imagined consequences of the second. Simondon’s third alternative, viewing technical objects as akin to living organisms (61), places humans squarely in the same ecosystem. With hindsight, I recognized that Q-Chords was always conceived as a tool, no more, no less. And as a tool, it acted in predictable ways, and could not evolve. The wall that my cultural background had built around my scores was mimicked in the minimizing of interactive possibilities for Q-Chords’ users. Yet, as Donna Haraway stated in her “Cyborg Manifesto”, there can be “pleasure in the confusion of boundaries” between human and machine. (Haraway 2003, 292) She acknowledges that pre-cybernetic machines seemed to their designers and users as not being “self-moving, self-designing, [or] autonomous.” (294) To me, the fixedness of Q-Chords made it a successor to Haraway’s pre-cybernetic machine. Though it moved, it had no animus, and it carried out its task slavishly and without resistance.

After I recognized this limitation, I understood that any systems I designed to produce music would need to exist inside an environment where they could be interactors themselves, encountering others, reacting and responding. To accomplish this, I would need to allow these systems the capacity to act independently of me in some manner that admitted the agencies of others into the unfolding of their sonic fabric. To do so, however, was to choose to give up some of my authority as a composer, and to create the conditions under which other voices might rise above my own, and take the sounds that followed into directions that I could not anticipate, and might not appreciate. There was a comfort in Q-Chords, one that arose from my knowledge that no-one could create conditions within the bounds of its parameters that would exceed my own intent. It was the same for my musical scores—my expression was locked within them, and I admit that at times when a conductor asks me for advice regarding how to achieve my intent in performance, even now I experience a small flash of self-satisfaction that all in the concert hall may defer to my will (even if they grumble about it outside of earshot). Now, I felt, I had to open up my music to a multitude of voices, and, as a generally non-performing composer, this troubled me. Could I create something as meaningful and engaging as my best linear scores, but working in a dynamic milieu? And would I disappear within it? This was a crisis of agency, which I will detail in the following chapter.

Chapter 6

Middle Works: A Crisis of Agency

Honour thy error as a hidden intention.

(Brian Eno, in “Oblique Strategies” (Eno and Schmidt 1975))

6.1 Introduction

After designing Q-Chords, I worried that I had not created a living, dynamic composing system, but rather a system that prioritized my agency as if I was still writing a linear score in Cubase. I had maintained my control over generated music through code—though the harmonic patterns produced by the program were variable, the code ensured that the output conformed to my objectives. My implementation of Q-Learning failed to draw upon an awareness of the patch’s users or its environment, for once it was engaged, its computations, though not strictly deterministic paid no attention to anything other than the dictates of the code. Its execution was frozen, and thus bore more similarity to a notated score than the dynamic music of a video game.

One value of Q-Chords was that it prompted me to question the value of relaxing my control over notation, adopting a more flexible approach to composing through code. There are many successful dynamic musical works, in video game applications and otherwise. Brian Eno’s Bloom, for example, specifies most musical parameters, but also grants its users an interface to perform the app. Lewis’ Voyager acts as an improvising partner for its musician users, creating a human/machine dialogue (George E. Lewis 2000, 33). In these applications, the programmer is not the origin of the experience, like Christopher Small’s dead composer (Small 1998, 92) or

Barthes' "Author-God." (Barthes 1977) The programmer is an origin, a participant, an observer, and a critic, in a dance of agencies where her agency appears not in dictation but in dialogue. In order to compose through code, I had to step down from my imaginary tower, and become willing to play with others.

Through developing the works described in this chapter I confronted my previously-unarticulated desire to maintaining agency over music structure, and resolved to change my understanding of my role in music. I feared, at the time, that I might disappear in my new musical medium, simply by relinquishing note-by-note control of unfolding events. To lose control thus seemed like the death of my agency in the service of others whose combined cross-talking would surely form a cacophony. But Barthes still wrote after he proclaimed the Author's death, his hand did not contract in rigour, and he continued to enjoy the act of writing. Discovering why Barthes continued to write could be a solution to my particular situation, which I describe as my "crisis of agency."

The projects documented here helped me to articulate the sources of this crisis. I realized that I had tacitly believed that the composer's unifying voice in music—which I perceived as the origin of the cohesive vision I assumed had to exist for a composition to be meaningful—was threatened by the integration of the capabilities offered by coding environments. Interactivity, randomization, sensor data input, and autonomy all seemed set to distance me from the music produced. While I did not know it at the time, I had inadvertently conflated the concept of *werktreue* with my understanding of the coding composer. *Werktreue*, as discussed earlier, is the dictate that a musical work is "held to be a composer's unique, objectified expression, a public and permanently existing artifact.... [it is] fixed with respect, at least, to the properties indicated in the score and it is repeatable in performances" (Goehr 1989, 55). I designed Q-chords to function as I intended, towards goals I determined, producing output I desired. It was *werktreue* in code, unchanging in its core, and in order to evolve I had to dismiss it, and embrace an alternative conception of my role and my work. The three projects described in this chapter are the trace of my struggles towards carrying out that process.

6.2 List of Works

The following works are described in this chapter, in chronological sequence:

- “Moonloops,” an interactive sound app realizing an ambient-styled soundscape.
- “Finger Filter,” a performance tool designed to extend and increase the accessibility of performance techniques for midi keyboard controllers.
- “Chief Paskwa’s Pictograph,” a dynamic, interactive soundscape designed for a museum in Regina, Saskatchewan, in support of an exhibit presenting a history concerning certain interactions between local Indigenous communities and English colonial settlers.

6.3 Moonloops

Moonloops was developed on a MacBook Pro laptop (2013 model) running MacOS. It was coded in the Max programming environment. As it was the first project I developed following Q-Chords, I intended the project to address some aspects of Q-Chords that I felt were problematic. These aspects included:

- Its manner of segregating user interaction and music generation processes, which occurred sequentially, rather than simultaneously.
- The simplicity of its musical output.
- Its narrow scope, in that it had only one primary function.
- The fact that I felt its contribution was almost entirely of a technical nature—I felt that it was best described as a tool, and not as an artwork.

6.3.1 Technology and Function

Moonloops was designed in Max as a modular system. The following descriptions of the modules reveal how the app functions.

Module 1: Event Generation

Module 1 generates note events (a tuple containing values for pitch, loudness, and duration) with randomized pitches. There are twelve possible pitch values representing the full pitch set of the chromatic scale. All pitches are tuned to equal-tempered frequency ratios with A set to the 440Hz standard. There is no weighting in the randomization process: each of the twelve pitch chromas are equally likely to be assigned to a generated note.

Note generation is triggered through an instance of a Max “metro” object which sends an impulse, known in Max as a “bang,” out of its outlet every five hundred milliseconds. A “counter” object receives each bang as an input, responding immediately by outputting an integer in a looping series ascending from zero to eleven. This number is routed into a “sel” (for “select”) object, which outputs bangs from its twelve outlets, each of which is activated by a different value from the counter. The select object’s outputs are routed to other objects that define and filter note pitches. Because of the regularity of the metro object’s bangs, it was possible to use the select object as a metronome, a rhythm generator, and as a means for enforcing timing sync between the various channels. Changing the rate of bang production is equivalent to altering the tempo of the performance.

Module 2: Pitch Filtering

After generation, all notes are passed through a pitch filter which accepts and rejects notes based on whether the pitches conform to a ruleset. Notes with conforming pitches were passed through the filter, rejected notes with non-conforming pitches were transposed either up or down to the nearest allowable pitch, and then passed through. Where notes are equidistant from two conforming pitches, the filter directs those pitches to the pitch that has the highest position in the tonal hierarchy for that key. The concept of the tonal hierarchy holds that a “a musical context establishes a hierarchy of tones” where certain tones are more “prominent, stable, and structurally significant than others” (Krumhansl and Cuddy 2010, 52). In pitch hierarchies emerging from tonal music in a major key, for example, the highest position in a pitch hierarchy is the tonic, followed by the third and fifth scale degrees, then the remaining pitches of the scale, and finally the pitches excluded from the

scale (Krumhansl and Kessler 1982, 336). There is evidence that these hierarchies also exist in human cognitive processes (Krumhansl and Cuddy 2010; Krumhansl and Shepard 1979), and in its filter module, Moonloops attempts to engaged this cognitive hierarchy, evoking a particular scale by effecting a pitch distribution consistent with its hierarchy.

Five rulesets can be used to configure the filter, the specification of which is under the real-time control of the user. Each ruleset is based on a scale known in tonal music. The first is the chromatic scale, which essentially acts as an open filter, since all generated notes have pitches which belong to this scale. the other scale options are major, harmonic minor, whole-tone, and pentatonic. I intended the scale-based filter to impact the emotional impact of generated themes. I will discuss the emotional qualities of the output of Moonloops in Section 6.3.3.

In addition to allowing the user to select from five filter modules that determine the scale, Moonloops includes a “channel transposer” function. Here, each of the five midi output channels can be transposed independently by the user by intervals of up to plus-or-minus six semitones. When a channel was transposed to a different root pitch from another, the sets of pitches passed through their note filters would move out of alignment, since the scale filter module precedes the channel transposer module. The consequence is richer harmonies, since by not aligning the root note of the tracks with one another, there is a possibility for harmonies to appear that are formed by pitch combinations not possible with certain scales. The whole-tone scale, for example, excludes every second pitch as the chromatic scale is ascended. By setting the filter to its whole-tone mode, and then transposing one track up by one semitone compared to the others, the full set of chromatic pitches becomes accessible, although harmonies that include semitone intervals can only appear when the transposed track is sounding at the same time as a non-transposed track.

The this module also includes a function that reduces dissonant note collisions, labelled “Inter-Voice Pitch Conformation.” Another instance of my lingering desire to exert compositional control over all app expression, this module acted as an additional filter. It monitored the pitches of all currently-sounding notes, across all five output channels, and, when a note was passed through the scale filter with a pitch that came within one, two, or three semitones of another currently-sounding

pitch, this filter would transpose it to the same pitch as the current pitch. The precise interval which triggers this effect is user-controlled, and because it can prevent the appearance of pitch combinations involving the semitone, which might otherwise be a common, dissonant occurrence in certain configurations of Moonloops the degrees of consonance and dissonance of the generated music was, to a certain extent, controllable in realtime.

Module 3: Tracks and Output

Moonloops is multi-timbral, with its filters processing five separate midi output channels, all routed via virtual midi cables into an instance of Ableton Live, a digital audio workstation application running concurrently on the MacBook. Instrumentation is predetermined and cannot be altered by the user. I chose sounds from the set of my available VSTi plugins, which are hosted in Ableton and controlled by outputs routed into Ableton from Max's virtual midi output ports.

During development I experimented with instrumentation provided by software synthesizers like Spectrasonics' Omnisphere and Arturia's Analog Lab, as well as Native Instruments' sampler plugin Kontakt, in which I hosted Orchestral Tools' Orchestral Grand Pianos sample library. Since Ableton receives only midi-compatible note data from Max, timbral choices are not determined in the Max scripts themselves, and during development I was most concerned with choosing timbres that subjectively complemented one another, a reflection of my still-present desire to maintain control of expression.

6.3.2 Creation

I gave Moonloops that title because the initial timbres I used gave the performance a hollow, otherworldly quality. I imagined Moonloops scoring a film scene following a lost astronaut lost deep in space, gazing upon stars that move with a glacial slowness. I also was inspired by the music of minimalist composers who have influenced me, such as Philip Glass or Steve Reich. Their music often feels suspended to me, still even when it is at its most energetic. Though its vertical, harmonic structure is variable through the scale and pitch conformance functions, its horizontal progression in time forms a somewhat static texture, like sections of Reichs' piece "Different Trains,"

for example. Its character also recalls the movements “Opening” and “Closing” in Philip Glass’s suite “Glassworks,” two arrangements of the same minimalist composition, the former scored for piano alone, and the latter for a chamber ensemble. In these compositions, Glass builds repetitive textures through instrumental lines that counter one another, twos against threes, producing beautiful syncopations. To me, Moonloops is like a segment of such a composition, extending an idea of a mere bar’s length into an endless meditation.

Moonloops’ development was also influenced¹ by Eno’s Bloom and Scape, introduced in Chapter 2. In their design Eno arguably raised the spirit of his ambient music aesthetic through his configuration of the app’s generative engine. I, too, have been intrigued by the idea of coding a simulation of my own compositional sound. Such compositional simulations have been designed before, including David Cope’s *Experiments in Musical Intelligence*, also introduced in Chapter 2. Unlike Cope’s system, from which the style of its musical output derives from the system’s analysis of a library of input scores by a particular composer (Cope 1992, 69-70), my aim was to design a system for which musical style emerged from coding choices. Again, this reflects my persisting, unspoken desire to create an mirror of my own compositional agency.

6.3.3 Reflection

The design process for Moonloops mirrored aspects of my compositional process in Cubase. This process could be described as improvisational, in that I would code in tandem with listening to outputs, as use the effects to guide my refinement of the three modules. The “Inter-Voice Pitch Conformation” submodule, for example, was not anticipated when I began the work, instead emerging after I encountered dissonance when differentially transposing multiple midi channels. Only after discovering the semitone collisions that ensued, and deeming them unpleasant, did I conceive of that submodule. This act was an early instance of comprovisation in a coding environment, a point-of-contact between my compositional and coding practices which helps to

1. Besides Eno’s work, Moonloops was also influenced by the textures of linear music such as the meditative music of the Icelandic band Sigur Rós, the wash of arpeggiated notes in certain minimalist works such as the track “Facades” from Philip Glass’ album “Glassworks,” and the background textures that I and others typically compose for film settings where unobtrusive musical support is a composer’s primary objective.

establish my recognition of a bridge between composer’s agency as projected through a score and as projected through coding.

The production of Moonloops accompanied a significant shift in my approach to creating dynamic music. In Q-Chords, I had a clear objective in mind: the automated construction of chord sequences modulating between two user-chosen major chords. The computer “performed” its script in accordance with the code I designed, and its generative processes were closed to input after they were engaged. A user might just as well been pressing the “play” button on a CD-player, for they had a similar degree of control over the form of the music once they clicked the “Generate” button. In Moonloops, by contrast, the interactor was not limited to the role of an initiator. The interface allowed for shifts in key, tempo, dissonance, and interactions between different instruments. It was an unfolding process that invited a degree of experimentation and play, both during its design and when a user interacted with it. If Q-Chords was a solution to a problem, Moonloops was an invitation to play, and had no defined endpoint other than the user’s choice to disengage.

That said, the code that animated Moonloops itself was quite fixed, still firmly subject to my coder’s agency. It had no memory, no model of its user, and it could not surprise a user, once they had worked out the scope of its controls. Thus I had yet to truly unfreeze my agency, a process that would require me to embrace the agencies of the users, and accept that I did not need to cling to a position at the very centre of music.

6.4 Finger Filter

As the name might suggest, “Finger Filter” was a Max patch designed to filter out undesirable actions performed upon a midi keyboard. I envisioned it as a means to make keyboard performance accessible to those who lack experience with keyboard performance, or whose fine motor control was otherwise impaired. Finger Filter would mediate between a digital keyboardist and a midi-controlled sound generator, and, as in Moonloops, reject playback of notes with pitches that failed to pass through a filter. At the same time, Finger Filter was designed to preserve some of the expressive possibilities of keyboard performance, while lowering the skill requirement needed to achieve a subjectively-agreeable performance. Paradoxically, I found that limiting

available pitch classes revealed new approaches to keyboard performance, extending my options for self expression when performing while mediated by Finger Filter.

6.4.1 Technologies

Finger Filter was designed in Max on a MacBook Pro (2013) running MacOS connected via USB to two-octave mini keyboard controller. Ableton live hosted the Kontakt plugin providing the piano samples, similar to the configuration in Moonloops. The script would intercept all midi note-on events and either pass them through to Ableton unchanged, or transpose their pitch and then pass them through, in a module essentially identical to the equivalent module in Moonloops. The user could the scale type—major, harmonic minor, whole-tone, or pentatonic—as well as define the root pitch.

Finger Filter further embellishes users' keyboard performances. It has an “echo” function, which could, when a key was pressed, generate two echoing tones, the first delayed by one second and the second delayed by a further half second. These echo notes were also passed through the scale filter, ensuring that they conformed to the user's chosen tonality. This setting added a textural quality, a familiar production technique often used in electronica, including We, The Artificial. Echoes conformed to three user-selectable motifs. In the first, a note performed by the user was followed by notes a fifth, and then an octave, above the played note. In the second, the played note was followed first by a note pitched a minor third down, and then a note pitched a minor third up, both relative to the original pitch. The third produced random pitches, drawn from the set of all pitches within an octave of the performed note. The random mode could be frozen, so that the randomly-generated intervals of the echoing notes could be consistent. All echoed notes were also passed through pitch filter.

6.4.2 Creation

Finger Filter was developed in 2014 while I attended the three-week-long Summer Institute for Critical Studies in Improvisation, in Saint John's, Newfoundland. The Institute included a course component in which I participated but was not enrolled. The course brought students from several countries together, many of whom were

talented instrumental performers. Because improvisation was a central theme of the course, there were ample opportunities for students to jam with one another, in and out of class time. I realized early on that my own keyboard skills were not on par with those of many of the other students there. I wondered if I could design a software tool that would enable me to perform alongside the other students, overcoming the limitations of my dexterity by creating conditions where notes I deemed undesirable would be filtered out. Thus began the brief design process of Finger Filter.

6.4.3 Reflection

Finger Filter is admittedly a fairly simple program, trivial in its design as similar filtering functions appear in numerous apps², including Bloom and Cubase. Its inclusion here is not meant to justify its innovation but rather how it shaped my view of the impact of coding on the agency of the musician: I had struggled to perform on an instrument, I added a simple computational mediator between us, and as a result it changed and extended my performance capacity. In one approach, it turned the keyboard into a percussion instrument where slamming the keyboard (an established keyboard technique) would produce a harmonic cluster that was surprisingly consonant when the filter was set to a pentatonic scale. In another, a glissando on the filtered keyboard offered a wider range of effects than on an unfiltered one, for the glissando could conform to scales (e.g. whole-tone or pentatonic) that could not easily be achieved by drawing a finger across an unfiltered piano keyboard. It permitted use of these techniques without altering or constraining velocity (volume) or timing of notes, retaining standard means of varying them expressively.

The paradox of Finger Filter is that, as it restricts the range of pitches that may be produced, it extends the utility of certain keyboard performance techniques, and lowers the degree of skill required to perform certain types of music on a keyboard. For example, engaging the filter prior to performing a chromatically-rich piano piece will

2. Finger Filter could be seen as a tool for performers, and similar technologies have been developed to help less-experienced composers to make their harmonies more complex by acting as a transformer of MIDI note data that evokes harmonies that the user may not have otherwise considered. An example is Xfer Records' MIDI plugin "Cthulhu," a combination of an arpeggiator and chord filter that can intercede between a sequencer's midi output and an instrument to transform the chord being played based on a progression of chords. While Cthulhu is aimed at digital composers, Finger Filter is aimed for performers, acting in real time to transform music as played. This is a small technical difference, but it means that the two apps affect their users in very different ways.

render it effectively unplayable if the objective is to uphold the dictates of *Werktreue*, but should a performance be attempted, it may still be expressive, albeit in a different manner.

While the innovation in Finger Filter is arguably slight in the broader context of musical interface design, it is included here as a marker of a personal shift in my perception of my relationship with machines. This perception began with my use of digital audio workstations—I appreciate the transparency of Cubase, recognizing that it provided a set of tools that helped me to realize *We, The Artificial*, at the cost of acknowledging that the album then became a product not only of my agency but the agencies introduced by those tools as well. In designing *Q-Chords* and *Moonloops*, on the other hand, I was greedy with my agency, fearful that if I were to give up too much of my precious control, the work would suffer as the divergent, conflicting goals of the independent agencies would descend to cacophony. *We, The Artificial* is unified in that all of its songs, in their instrumentation, the subject of their lyrics, the vocal performances, and indeed every feature of the song is rooted in my choice, a product primarily of my agency. To embrace the interactive score, the musicking machine, I needed to find myself within the code, not as one voice hidden among many, but rather a voice that sang loudly, and, crucially, with a magnified expression. This did not occur in *Q-Chords*, I could have composed better transitions with more ease in Cubase. It did not occur in *Moonloops*, because, even though I included an interface to alter the harmonies, the work in full didn't surprise me, and its capacity for expression was limited not by what you did with it, but by the form of its code. *Moonloops* never really surprised me. *Finger Filter*, however, did.

By acting as an intermediary between my keyboard performance and the output music, *Finger Filter* invited me to explore keyboard techniques that I wouldn't

otherwise have used³ . Using the simple script, I discovered a mode of expression—realized through extending the set of keyboard techniques available to me—that was new, and surprised me. I enjoyed it immensely, for in an instant I became a new kind of keyboard performer, or a performer on a new kind of keyboard, and discovered a magnification of my ability to engage in self-expression. This was the breakpoint that I had been seeking, the point where my coding magnified my voice, instead of forming a hollow shadow of it, or of drowning it out in a sea of others’ voices. Finger Filter was the first instance of what I consider a magnification of agency in the chronological series of coding projects I undertook during my doctoral program, and so marks the beginning of my identification as a composer whose agency is projected through acts of coding.

6.5 *Chief Paskwa’s Pictograph*

The final project in what I understand to be the transitional period marking my shift from assuming the role of composer through notation to assuming the role of a composer through code is titled “Chief Paskwa’s Pictograph” (abbreviated “CPP”). It is an interactive soundscape designed to be installed alongside a museum exhibit in Regina, Saskatchewan. Due to events beyond the scope of the interactive music project, the exhibit was ultimately not installed. The experience, however, allowed me to explore the limits of my agency with respect to honouring voices that are not my own, and it was a formative project that affected how I approached my final project, “Shards of Memory,” particularly in its methods of mingling agencies.

The creation of the interactive score was funded through the NSERC Engage program, and was developed through a partnership between the University of Regina and Computer Science Professor Howard Hamilton, Twisted Pair Sound Studios, and

3. In a sense, Finger Filter shares some of the essential functionality of an EDM DJ rig, which can allow such DJs to engage in real-time performance, triggering prerecorded samples or audio effects, potentially through a keyboard controller interface. Such a setup appears similar to the configuration of finger filter, in that keypresses in both systems trigger musical effects as part of a live performance. That said, in Finger Filter, the functions triggered by the performer’s keypresses are focussed towards keyboard performance, and leverage the keyboardist’s instrumental skill to achieve expressive ends. A DJ need not conceive her keyboard controller as anything other than a series of buttons. A user of Finger Filter, when treating the keyboard as a piano or organ, may find new expressive possibilities for keyboard performance. The apparent similarity between Finger Filter and a DJ rig belies the former’s expressive potential.

the Royal Saskatchewan Museum.

6.5.1 To Integrate Sound Into a Museum

What is the role of sound in a museum? According to Harald Kraemer, the question as to how multimedia—including sound—might be incorporated into the museum space has been of great concern within these institutions since at least the 1990s (Kraemer 2018, 82). That, he observes, was the time when media applications such as Laserdiscs and CD-ROMs became a primary platform for storing and sharing knowledge (83). These objects carried a *Gesamtkunstwerk*, like the band Kraftwerk introduced in Chapter 4, in which multiple media, such as sound, animation, and text, combined to form a then-novel experience, a “total work of art” that generated excitement among the public (83). Conceiving of a museum in such terms, however, can be cause for concern. Some have feared that the museum would turn into a “Disneyland,” where enticing visitors to engage with spectacles—games, technological novelties, plays of light and sound—would distract them, reducing the effectiveness of the exhibits by obscuring the messages they are intended to project (Ruttkey and Benyei 2018, 105). This is, I feel, a worst-case scenario that results not from a flaw in conceiving of the exhibit as a *Gesamtkunstwerk*, but simply poor design. That said, given the subject matter which my soundscape was intended to support, I was wary of allowing the soundscape I was developing to distract patrons. This trepidation, of creating a spectacle that obscured the important message of the exhibit, became the force that challenged my developing theory of composers’ agency, and made me rethink my approach to composing through code.

CPP was an interactive, immersive, site-specific soundscape intended to engage and support the education of museum patrons as they traversed the exhibit. I wished to foster an empathetic understanding among patrons towards the Indigenous people whose historical narratives formed the subject of the exhibit. There was also a technical problem motivating me: I had observed that museums typically integrate minimal audio content into their exhibits, and wanted to use interactive technologies to change that. I reasoned that the underdeveloped sonic aspect of museums was prolonged by concerns that a buildup of unrelated sounds emanating from multiple displays would form a cacophony that would distract patrons from the ideas the displays are designed

to convey. It occurred to me that this issue is not unique to museums. Video game sound designers frequently face situations where they have multiple sound sources, spatially arrayed in the virtual game world, which must be mixed in such a way that the most important sounds are audible and understandable to the player. I realized that if I approached the exhibit as if it were a game environment, I could design a soundtrack incorporating multiple sound sources governed by a central script which would mix the sound sources in real time. If the system had sensors that could identify where patrons were, it could prioritize aspects of the sound mix associated with their location in the exhibit, and at times when the exhibit was particularly busy, the central script could attenuate certain sound elements to avoid the cacophony that would otherwise result. This approach, drawing principles of video game scoring out of the game and into the real world, could theoretically lead to a music of museums.

With that feature-set in mind, I began to build the system.

6.5.2 Form of the Installation

The path through the exhibit traced a chronologically-arranged sequence of six historical displays, as follows:

1. *Intro/Entrance*. The immediate entrance of the exhibit, consisting of an arch and welcoming text.
2. *Pre-Contact*. A depiction of precolonial life in what would become the province of Saskatchewan.
3. *Treaty 4*. An exploration of the nature of treaties between Indigenous people and colonizing settlers, providing context for a discussion of the history and meaning of Treaty 4.
4. *The Pictograph*. An account of the conditions under which Chief Paskwa, of the Cree First Nation, drew a pictograph depicting items that were owed but not provided under the terms of Treaty 4.
5. *Repatriation*. A history of the Pictograph document, which travelled to England and spent more than a century there, before being repatriated to the Paskwa First Nation in 2007.

6. *The Future of The Pictograph*. The actual Pictograph is displayed, alongside written reflections on the text provided by local Elders.

As can be observed, the design of the exhibit space allowed the patron's path to reveal a narrative of oppression and a call to preserve history. Perhaps ironically, that history is presented in a museum that is an instance of an institution that has a troubling record relating to the presentation of Indigenous histories. By contributing to the exhibit I was effectively entering that space where history meets colonial narratives. As such, I will now introduce some commentary on how I approached this issue.

Questioning the Place of My Agency

The exhibit explored the history and consequences of nineteenth-century treaty negotiations involving First Nations and the colonial Canadian Government. These treaties have not expired, and the site of the museum remains treaty land. Treaties are binding legal agreements between First Nations and Canada's government, which remains controlled largely by people of Colonial descent. In his book "The Inconvenient Indian," Thomas King (a Canadian-American writer of Cherokee and Greek ancestry) examines the relationships between Indigenous people and settlers as part of an exploration of Indigenous identity. He writes:

A great many people in North America believe that Canada and the United States, in a moment of inexplicable generosity, gave treaty rights to Native people as a gift. Of course, anyone familiar with the history of Indians in North America knows that Native people paid for every treaty right, and in some cases, paid more than once. The idea that either country gave First Nations something for free is horseshit. (King 2012, KL.1001)

In this light, the development of this installation required sensitivity to the fact that it was supporting the telling of an ongoing history of colonial oppression, one that my own cultural biases might lead me to mischaracterize.

Some aspects of the exhibit's narrative about Treaty 4 and the Pictograph explore injustice, broken promises, and the silencing of Indigenous voices. These issues continue today. At the same time, the institution of the museum has had a complex and

at times deeply problematic relationship with Indigenous people and the recounting their histories. A current example of such a controversy relates to the Canadian Museum of Human Rights. In a article published on July 12, 2020, the CBC reported⁴ that Black and Indigenous employees accused the museum of “racism, discrimination, and censorship.” These claims are part of a longer history in which the article identifies concerns of political interference and negligence with regard to the preservation of artifacts. Underlying these concerns is a recognition that the museum can become a site that does not transparently transmit Indigenous voices or histories, as explained by Susan D. Dion and Angela Salamanca:

Anthropology museums are concerned with collecting, salvaging, and exhibiting objects representing material culture of what they deem to be authentic and vanishing Indigenous people. Art galleries represent the history of “aesthetic” objects by Western peoples from a distinctly western perspective. Contemporary Indigenous artists don’t see themselves in the museum and, until recently, have been marginalized by the gallery. Acknowledging Indigenous experiences, perspectives, and histories challenges the dominant narrative.... (Dion and Salamanca 2014, 165)

My doctoral studies have centred on the question as to how I, as a composer, might turn to programming as a compositional practice, and how the transformation of my methods for exerting my agency might benefit were I to contact other agents through coding algorithmic music systems. While I had, to this point in my studies, held on tightly to my capacity for control over music generated by apps I designed, I recognized that doing so in this project might obscure the histories which it purported to present accurately. The exertion of my agency can bring with it my experience, biases, values, and understanding of history, and this is good where I attempt to make a work that expresses my perspective. However, in the context of an audio installation intended to support the exploration of Indigenous histories, it was difficult to see how my perspective was relevant.

I did not wish to create an interactive sound world where the music, sound effects, and interaction mechanics changed the meaning of that history—to do so, I felt, would

4. The article, by Kelly Geraldine Malone can be viewed here: <https://www.cbc.ca/news/canada/manitoba/canadian-museum-human-rights-winnipeg-history-controversy-criticism-1.5646802>

risk having the history become obscured or colonized. It was a situation unlike a film score, where my expression is typically free to range within the broad parameters that the film's director ascribes. That is a position of power, where my agency is invited to engage in self-expression. But the sound world of Chief Paskwa's Pictograph was not supporting a history I could fairly interpret. Using a chord pattern that evoked particular emotions inappropriately or triggering sound effects that detract from the telling of history would constitute an erroneous interpretation, and do damage to the authentic history. Those kinds of expressive choices, which I have made freely as a film composer and in my earlier algorithmic music works, did not belong here. The sound world was not about me, it was about a history that I, as a white man of colonial descent, might struggle to represent accurately, and lack a frame of reference. For this reason, I chose to attempt to retreat from control of the installation. I consciously curbed my agency, deciding to minimize the use of my generative music, and focused on collaborating with Indigenous musicians and using authentic recordings of locally-sourced sound effects.

This choice was problematic in that, by retreating from fully embracing the application of my agency, I removed from my process some of the skills and self-expressive depth that might otherwise have brightened the project. That is not to say that my decision was incorrect—absolutely, I did not want to further colonize the space by filling it with my own expression at the expense of others'. This choice, while I maintain it was necessary, caused me to question something that I have not considered: throughout my doctoral studies, I had been trying to avoid prioritizing my agency, to better make use of the code medium in which I was beginning to compose. I had always felt that I was too privileged in projects before CPP, and this is part of the reason why my instinct in starting CPP was to withdraw. In this act, however, I discovered that withdrawing from my interstice algorithmic music systems could also detract from the strength of these systems. I made the best possible choice to remove myself as much as I could, and was rewarded with the realization that my past bias, towards seeing interactive music projects in terms of an opposition between creator and user. This project, by forming the mirror image of the unbalanced agency I saw in my previous projects, allowed me to realize that I needed to frame my algorithmic music as a site of shared agency, and worry less about the distinctions between roles.

The History of Chief Paskwa's Pictograph

The Pictograph itself was drawn by Chief Paskwa, a using iconographical depictions of goods that were promised to his people under the terms of Treaty 4 (Beal 2007, 110). That Treaty was negotiated by Lieutenant-Governor Alexander Morris of Manitoba and the North-West Territories with the Cree and Saulteaux Nations in 1874, though it is notable that Nations were present in the area covered by the Treaty, such as the Assiniboine (116). The treaty covered 195,000 square kilometres of territory, including much of Southern Saskatchewan and parts of Southeastern Alberta and west-central Manitoba (Stonechild 2005). The Nations agreed to Treaty 4 on the basis that they would be paid annuities such as tools and farm implements, reserve land, annuities for chiefs, and a Treaty Ground to serve as a meeting place to receive promised goods (Stonechild 2005; Beal 2007, 120). However, provision of these goods was slow, and the Treaty Ground was ultimately abolished, partly due to colonists' fears of large assemblies of Indigenous people (Stonechild 2005). These deprivations were compounded by the fact that agents of the Canadian Government failed to provide adequate food during a famine subsequent to Treaty 4's signing (Daschuk 2013, 143).

Chief Paskwa, a Plains Cree by birth who lead a Saulteaux band, was a signatory to Treaty 4 (Nestor 2005). His Pictograph, formed from his personal iconography and not from a form of Cree writing in , represent the Treaty, though the icons' precise meaning was only known to Chief Paskwa, and no-one appears to have discussed their meaning with him (Beal 2007, 122). The pictograph represents Chief Paskwa's personal record of the Treaty's most important points (122). The document has two panels: the Treaty panel, effectively a transcript of the negotiations (122), and the Provisions panel, thirteen lines of symbols representing treaty supplies promised and provided (130), showing items that had been withheld, such as two pigs and six ploughs that were owed to his people, but had not been supplied (141). In 1883, the document was provided to Henry Barnaby, an English traveller to the region. The reasons for this transfer are unclear, but may have something to do with Paskwa's wish to send a message to English authorities to request help in receiving the provisions promised (145). Instead, the document languished in a private collection in England for most of the years intervening between its creation and its repatriation

in 2007. As a consequence of the treaty’s terms being left unfulfilled, the Nation experienced famine during a particularly harsh winter, during which many members of the Nation needlessly died as the buffalo herds disappeared. The colonial government nevertheless, continued its policy of not honouring its treaties, causing much hardship (Beal 2007, 143).

As I was tasked to help share this history of colonial oppression, and in consideration of my identity as an English colonist living on Treaty 4 land, I resolved to be sensitive to the history and to try to present the history as faithfully as I could. Historically, research has been “completed on, rather than with (i.e., in collaboration with) Indigenous Peoples in Canada” (Drawson, Toombs, and Mushquash 2017, 1). I did not want my colonial identity to obscure or misrepresent any aspect of this history through poor design of the form or content of the the audio installation. This led to my choice to try to curb my agency in carrying out certain design processes. The choice proved to be significant and instructive, as I will explain in the following sections.

6.5.3 The Sound Engine And Its Technologies

The CPP sound installation made use of the following equipment:

- Apple Mac Mini computer running Max and Ableton Live. (1)
- Eight-channel firewire digital audio interface (with a Thunderbolt adapter) (1).
- Multichannel power amplifier (1).
- In wall speakers (6).
- In wall subwoofers (2).
- Request-to-exit motion sensors (6).
- Arduino programmable circuit board (1).

The entire soundscape was generated in realtime by a Max patch I designed. The Max patch was capable of multichannel audio output to an array of eight speakers, and of using sensor inputs to instantiate and modulate sounds produced by any of

these channels. I developed the installation in a sound studio at the University of Regina campus, rather than at the museum itself. Figure 6.1 contains photographs of the testing space where the soundscape was developed and refined. In this image, you can see how the speakers and motion sensors were arranged, in an approximation of the pathway which patrons would follow as they traversed the exhibit.

Sensor Inputs

The Mac Mini received input from the six motion sensors, each of which was installed in a specific location and facing a known direction within the exhibit. The motion sensors were able to transmit a value that was either “on” when there was movement detected in the sensors field-of-view, or “off” when movement had not been detected during the preceding five seconds. Altogether, the field-of-view of the six motion sensors covered most of the exhibit space, and care was taken to avoid overlaps between the fields-of-view of individual sensors. The sensors were connected to an Arduino circuit board, with each sensor connected to a different input on the Arduino. The Arduino was connected via USB to the Mac Mini.

An Arduino is a small, programmable microcontroller that features multiple inputs and outputs that can be monitored and controlled by uploading software to its chip. Arduino are commonly used for building interactive systems⁵. I installed a library supporting the Firmata protocol⁶ onto the Arduino, allowing it to receive sensor outputs and transmit their status to the Mac Mini. The Firmata installation configured the Arduino to appear to the Mac Mini as a midi port. This port was connected to Max’s midi input, allowing Max to receive sensors’ status as simple midi messages.

Audio Outputs

The installation included eight speakers: six mid-range speakers, and two subwoofers. All were to be mounted in the walls of the exhibit, with the mid-range speakers placed high in the walls to enhance audibility, and the subwoofers installed near the floor. Sonic output from the mid-range speakers was localized, and the mono

5. The official Arduino website is found at <https://www.arduino.cc>

6. More information about Firmata and its use can be found at <https://github.com/firmata/protocol>



Figure 6.1: Photographs of the development space where I designed and tested my audio installation for the Chief Paskwa's Pictograph Exhibit. The room is a sound studio in the Education Building of the University of Regina Campus.

audio channel for each speaker could be associated with sound material relevant to the nearest display of the exhibit. The two subwoofers projected bass frequencies for all six output channels, and acted like an acoustical “glue” for the localized sounds.

Each of the eight speakers was connected using speaker wire to the eight-channel amp, which was in turn connected via unbalanced 1/8” audio cables to the eight channel firewire digital audio interface mounted on the same rack. The digital audio interface connected directly to the Mac Mini computer via a firewire-to-thunderbolt convertor, appearing to Max as eight separate mono audio outputs.

Triggering Interactive Sound

Realtime behaviour of the soundscape was governed by two factors. The first was the coded function of the Max patch, and the second was the motion sensor inputs. Activation of the motion sensors was required to trigger sound generation, since after a period of three minutes during which no motion sensors were activated it was assumed that the exhibit was empty, and the space would go silent.

Since the sensors were situated at known locations throughout the exhibit, any sensor signal indicating motion could be pinpointed to a specific region of the exhibit. When a motion sensor was triggered it would be detected by the Max patch, which in turn would activate the soundscape elements associated with that sensor’s location in the exhibit.

If only one person were traversing the exhibit from entrance to exit (or if multiple people were traversing it but were staying close together), the soundscape elements would trigger sequentially, like a film score supporting a temporally-sequenced narrative, albeit with indeterminate times spent on each of its six scenes. In the more interesting case where multiple motion sensors were activated simultaneously, multiple patrons are assumed to be spread through the exhibit. In this case, where multiple audio channels are to be engaged, the Max patch begins to prioritize audio channels, based upon decision rules that I designed, and which as a consequence embody my agency.

6.5.4 Audio Content

The content of the soundscape is sonically diverse, reflecting the fact that the exhibit features many differing scenes in its attempt to faithfully represent the history of Chief Paskwa’s people and their Pictograph. The overall soundscape includes playback of sampled instrumental sounds triggered by midi within the Mac Mini box, sound effects sourced from field recordings and sound libraries licensed to project partners, and recordings of performances by Indigenous musicians Teddy Bison (voice and drum) and Trenton Heywahey (traditional indigenous flutes).

The following paragraphs describe the audio materials and their use in CPP.

Stage 1: Ambient prairie soundtrack

When patrons entered the exhibit space, I felt that there should be a sonic element that invited patrons to engage while establishing the exhibit’s setting. I used natural sounds, such as the calls of local wildlife and sound effects such as rain, wind, and thunder. I felt that such sound elements would quickly and effectively establish the backdrop of the histories it supported.

The ambient prairie soundtrack was not merely a loop. It was a responsive to motion sensor triggers, and it featured seasonal modes: “winter”, grounded by the sound of a cold, biting wind, and “summer”, featuring the sounds of a summer breeze, distant thunder, or a light rain. Both modes included wildlife sounds, authentic to the local region, drawn from field recordings provided by project partners. When the season shifted, a flock of Canada geese could be heard, marking the change in season. Triggering the speakers in sequence, the geese appeared to be flying from one end of the exhibit to the other, with their direction determined by whether they were migrating north or south.

Seasonal shifts were triggered by the motion sensors. Motion sensors one (entrance) and two (precolonial prairie) would trigger the summer mode. Sensors three (the signing of Treaty 4) and four (the hardships of Indigenous people which led to the creation of the Pictograph) would trigger the winter. Given its association with a narrative of grief, famine, and oppression, the winter mode was stark, while the summer mode, which served as the introduction to the exhibit, was more inviting. Of course, when the exhibit was busy, signals from both pairs of motion sensors could be

occurring simultaneously. In such cases, the localized musical content was prioritized, and the ambient sounds were mixed down in volume or silenced entirely.

Stage 2: Flute and generated chordal accompaniment

This portion of the soundscape accompanied the precolonial region of the exhibit. It featured a chord sequence generated in Max triggered by motion sensor one, and Heywahey's flute performances triggered by motion sensor two. Heywahey was recorded at Twisted Pair Studios, improvising melodies according to local Indigenous traditions in a sound-dampened studio. The chordal sequence acted as an accompaniment for the flute melodies, and together these two elements evoked a haunting, but unobtrusive, atmosphere.

Stage 3: Songs

This portion of the sound installation was developed in collaboration with Teddy Bison, an Indigenous singer and drummer who performed traditional songs for use in the Treaty 4 portions of the exhibit. Bison performed songs on three themes: a Chief song, a Warrior song, and three Sadness songs. These songs, like Heywahey's performances, were recorded at Twisted Pair Studios early in the development of the soundscape. Since the performances were of traditional Indigenous songs, I chose to avoid triggering additional sound elements during playback of the performances, as these songs were meaningful, and to obscure them would reduce the expression of both music and lyrics.

Bison's Warrior and Chief song performances were triggered by Treaty 4 motion sensor, a decision informed by Bison's observation that such songs would be performed at gatherings like a Treaty signing. The three Sadness songs Bison performed were triggered by the motion sensor in the region of the exhibit introducing the Pictograph, and effectively reinforced the tragic narrative documented there.

Stage 4: England

The Pictograph's travels to and from England were a period of lost time during which the Pictograph languished in a private collection. I approached this section of the narrative from my perspective as a film composer, since at the time I considered

the depiction of England to be one for which I could speak. I sourced recordings that I felt would evoke a Victorian cityscape. Samples were triggered by motion sensor five, with a thirty-second lag before re-triggering was enabled.

Samples included horses hooves and carriages, sounds simulating activity in a Victorian market, clock tower chimes and church bells. These samples were triggered in a fixed sequence intended to evoke a carriage journey in a Victorian city. Museum patrons would, I hoped, feel like they had been transported out of Saskatchewan, into a different kind of space. This choice to refrain from mixing these sound elements with those of other exhibit regions was my response to the difficult history being represented. I felt that taking the Pictograph to England, but storing it instead of presenting it to Queen Victoria, was a silencing of a vital Indigenous call for justice. In that light, integrating the London material seamlessly into the overall soundscape would have implied a connection that was not actually there, propagating a harm to Indigenous people.

Stage 5: Pictograph display and exhibit conclusion

At the end of the exhibit, the Pictograph itself would have been displayed in a custom-built case. This concluding section of the exhibit was outfitted with a speaker and motion sensor, and the production team wanted to parter with an Indigenous musician or musicians who performed in contemporary style. We reached out to two Saskatchewan-based hip-hop artists to request original material, and also considered already-released recordings that could have been featured here, but production of the exhibit concluded before a partnership could have been formed.

6.5.5 Reflection: Problems of Agency

The choice to collaborate with Bison and Heywahey reflected my belief that I could not create an effective installation if I were not to work with Indigenous artists. The Indigenous narrative had to be treated with sensitivity, and could not solely be informed by my own interpretations of that history. I feared that my perspective as a colonist would, explicitly or implicitly, act to obscure or distort aspects of the historical account, which would constitute a harm. I chose to curb my agency, avoiding the integration of interactive music and sound elements not derived, at least in part,

from Bison and Heywahey's performances, and restricting other sound material to prerecorded sound effect, such as natural weather sounds and foley. The choice to curb my agency proved unsatisfactory, but in helping me to better understand my role in my practice of interactive music, served as a beginning of a more robust phase of creativity.

Speaking for Others in the Museum

In an article critiquing the "colonial politics of recognition" and its relationship with the practices of Canadian Museums, Kelsey R. Wrightson observes that casting museum spaces as "either reproducing colonial violence or as the emancipatory reclamation of history and self-recognition" are an oversimplification of the complexity of museum practice (Wrightson 2017, 36). She characterizes "recognition" as the "hegemonic expression of self-determination within the Indigenous rights movement in Canada," and associates its appearance with a shift towards legal and political recognition of Indigenous rights, including that granted by treaties (37). The "politics of recognition" typically concerns "the relationship between the institutional accommodation of collective difference, and the freedom and autonomy of individuals" (37). While the aim was just relations, such efforts often acted to sustain the status quo (38). The museum plays a role in forming such relations, as did my CPP installation, by extension.

In this Chapter so far, I have indicated how, while developing this project, I tried to dismiss my agency, and so not obscure the historical account the installation was intended to support. How do other museums work in the light of the delicate relationship between Indigenous histories and such institutions?

The institution of the museum has been problematic with regard to its relationships with and depiction of Indigenous people. For example, Calgary-based Glenbow Museum's 1988 exhibition of Indigenous artworks, "The Spirit Sings: Artistic Traditions of Canada's First Peoples," distinguished itself by having no Indigenous people among its six curators, and a lack of Indigenous-sourced commentary (Gibbons 1997, 310). Perhaps as a consequence, the exhibit was criticized for homogenizing the origins of artifacts on display, conflating diverse Nations and situating them as history (Wrightson 2017, 39). Furthermore, the exhibit was exclusively sponsored by Shell Oil (a company that had long been surveying and drilling in the disputed land of the

Lubicon Lake Band of Cree) led the Band to publicly denounce the exhibit (Wrightson 2017, 39), prompting a well-publicized boycott of the Glenbow (Gibbons 1997, 310). While the debacle led to a fruitful new dialogue between Indigenous leadership and representatives of the museum community (Wrightson 2017, 40-42), the episode highlights how a lack of fair, informed representation can damage a community and sew mistrust.

6.6 Finding My Voice among Many

I had originally intended to approach CPP as if it were a dynamic film score, creating musical themes and sound effects that reinforced the historical narratives associated with each region of the exhibit. A film composer, however, assumes a position of authority, since their skill set and interpretive sensibility grants them a distinctly privileged status⁷. In my view of the dynamic between composer and director, the composer has agency over the creation of the score, but the director or producer has agency over what gets included in the film. The collaboration involves an alignment of these agencies. In this light, to ask John Williams, or indeed any film composer, to compose a score without granting them freedom of choice could be seen as a denial of their agency, and their role.

In crafting CPP, I recognized that this model, where I acted as a film composer who sought to serve the historical narrative itself, could not work, because in its requisite alignment I would necessarily impose my values upon the telling of the history. I pulled back, and tried to let the sounds speak for themselves. This approach was less successful than I would have liked. The technologies functioned, sensors triggering sounds which filled the development space. Yet I felt it to be hollow, with my deliberate disengagement hobbling my composerly agency. Where in Finger Filter the code magnified my expression, here my relinquishing of my agency to disembodied sound and history made me timid. If I were to do it again, I would have to find myself in the exhibit, or acknowledge that I did not belong there at all. Thus working through

7. Consider the relationship between director Steven Spielberg and composer John Williams. Spielberg hired Williams to compose scores for his films almost exclusively, and attributes the cementing of their partnership to Williams score for “Jaws,” in which Williams ignored Spielberg’s temp music and to follow his own vision, contributing to the film’s success and resulting in Williams’ first Oscar. See <https://www.latimes.com/entertainment/la-xpm-2012-jan-08-la-ca-john-williams-20120108-story.html>

CPP was a valuable experience, for it revealed that the ghost in the machine might be me.

Upon reflecting on the experience of the installation in its demo form, I had to admit that I was not absent from the unfolding sounds—far from it. Walking through the demo space, one might immediately notice the overall consonance of the sound. There was a clear attempt to order placement of sound elements. Transitions between regions of the exhibit were smooth, not abrupt. The sound texture minimal, not aggressive. The interplay of flute themes localized to the second and third areas of the exhibit behaved like a transition between two scenes in a film score. These features were all reflections of my values, embedded through my design choices. Though I had tried to pull back, my agency was still at work, its imprint active in the form and behaviour of the sound elements. Yet I had placed not one note of inviolable pitch and timing into the mix that had been subject to my specification. It was evoked by the code, sustained by the rules governing its behaviour, despite my attempts to banish it from the project. My agency had simply passed to another medium, the medium of code, and I found my music there.

In the earlier Moonloops project, I attempted to support interactivity with static code that was governed by my personal aesthetic values. In informal conversation with Moonloops users, the criticism arose that their interactions with it had too little impact upon the overall sound quality. I took this to mean that the form of interactivity I had built into the Max patch had failed to involve them sufficiently. Yet I now recognize that such behaviour was actually implicitly by design, since I had actually approached Moonloops as if it were informed by *Werktreue*. With CPP I tried to minimize my self-expression, yielding to others' expressions, such as Trenton Heywahey and Teddy Bison, and indeed to the patrons triggering the sensors. I intended all actors within the exhibit space to musick together. But I had withheld my own spirit, its agency elaborated over decades of composing and collaborating with other musicians, film makers, and dancers. In both projects the issue could be framed as an ill-informed balancing of agencies.

The crisis of agency is not simply an example of my mistakenly taking too much or too little control. It is a problem with how I conceived and formed the web of social relationships around these works, the site where musicking occurs. Eno's "oblique strategy" quoted at the beginning of this chapter connects mistakes with

“hidden intent,” and I feel that some of the struggles I had producing these three works stemmed from my unarticulated choice to set myself apart from others who encounters and contributed to the works. My crisis, I gradually realized, could only be resolved by finding within my projects a balance between my self-expressive voice, as expressed through code and notation, and my desire to join with others’ agencies. Like Goldilocks, I sought a balance that would lead to satisfying musicking in an algorithmic context. Once this goal was articulated, I became more confident in designing my culminating doctoral project, *Shards Of Memory*, where I embraced the contributions of others to its form, while remaining true to my own identity.

Chapter 7

“Shards of Memory:” Augmenting Space to Augment Social Bonds

7.1 Introduction

I consider each of the original works introduced in this dissertation to be an attempt to refine and extend my conceptualization of my own agency as it relates to the composition of music through digital technology. Taken chronologically, the works presented so far reflect an emerging realization that my previous conception of myself as *composer*—specifically in the form I understood it, which is most compatible with Small’s characterization of the deceased composer of classical music—was problematic when applied to the task of creating interactive music. My framing of the composer, always sitting apart from musicians and her audiences, became an unconscious model framing my identity as new media artist. I have introduced Simon Waters’ term “performance ecosystem” as a place where music is a social practice and the application of computer technology introduces ambiguity in the supposed boundaries between human roles (Waters 2007, 2-4). This concept challenges my earlier understanding of what it is to be a composer.

I gradually realized that if my works were to be experienced as an inclusive act of musicking, I would need to enter the ecosystem, deliberately. I had to accept that my self-expression did not need to be the dominant species that devours all the others, and that meaning is something that can emerge from interactions between people, or from other agents’ selves. According to Foucault and Barthes, my musical self-expression cannot reach the others with whom I musick unchanged. I believe their contentions

that language is a filter for reality and that meaning is ultimately a construction of the reader’s mind, even though there is a argument to be made that the Author sets the process of writing in motion, and hence owns some of its projected meaning. But the elevation of the reader as the agent who reconstructs meaning arguably applies to music as well as the spoken word, or even more so, given its semantic ambiguity. My early doctoral works, where I created interactive sonic art of which I was the primary agent (“Q-Chords” and “Moonloops”) or from which I retreated (“Chief Paskwa’s Pictograph”), can be understood as attempts to govern meaning through situating my agency at the extremities of control over the construction of the work. Upon reflection, those experiences revealed to me in practice what Waters observed of the relationship between the recording and the computer—that the former helped to compartmentalize music according to the roles surrounding it, whereas the computer could blur them (Waters 2007, 2). My chief concern in crafting interactive music systems had been to control the music that emerged, in a manner that recalls the fixity of recording. Q-Chords was not interactive at all, once the “generate” button was clicked. The music of CPP had less to do with my personality, although through it I facilitated the interactive presentation of the music of others. I came to understand that I needed to think of my artworks as social, musicking experiences of which I was participant, and not concern myself so much with where precisely I appeared in the work, and how important that position was. To become part of a musicking web of agents would be a success in itself.

My goal in “Shards of Memory” is to form a place where people can encounter one another as participating agents, and so form new empathetic connections. See Appendix D for a link to a video demo of the app as it functions. This goal is reflected in the title of the work. I describe the treasure we each have—our being—as a gem, each facet revealing a modicum of our unique human experience. I describe the act of coding as the shaving off of a piece of that gem as a shard, and then entrusting it to the computer that animates it. The gem belongs to the coder, the but the shard does not, for once the coder decides to give it up to the code, the shard may act autonomously, and so attain a life beyond the coder. This form of sharing through code is, I believe, unlike any other form of sharing through technology. Non-algorithmic sharing through a notated score may be fixed, specifying action without a spark of awareness, forming a thoughtless inertia. A traditional notated score

cannot reorganize itself, a piano keyboard cannot play itself, a hammer has no will. But when governed by an algorithm, a score may be animated. This animation (to extend my analogy) happens when the coder shares a shard, by creating code that acts according to some model of the coder's values. In that act, I suggest that a spirit emerges, distinct from the coder. This artificial impetus, which I consider to be a spirit because it can act autonomously, is a gift of our spirit which we endow into code, to let it act beyond us.

“Shards of Memory” explores the nature of connections among humans that form through contacts between agents, both human and artificial. I use the creation and function of the app as a means to explore the idea that a meaningful interpersonal connection may be formed through contact mediated by the acts of coding and the projection of personal agency that manifests as a consequence.

7.2 Goals

The “Shards of Memory” project aimed to complete the following goals:

- To disrupt my practice as a composer of linear music, particularly cinematic symphonic scores, through my design of an autonomous, interactive, dynamic score that may be functionally distinguished from my past linear notated scores, while still projecting a recognizable extension of my composerly agency.
- To reverse the previously-observed retreat of my agency from that of others who interact with my works, and build social meanings through my choice to join in a mingling of many agencies present in the space of my work.
- To counter the dehumanizing or divisive capacity of social technologies (seen, for example, in acts of trolling, desensitizing participation in simulated violence, and the wilful spread of disinformation), by creating a work that leveraged aspects of such technologies to promote empathy and increase social bonds between interactors who represent themselves authentically.

These goals respond to the issues that I identified in the earlier artworks documented in Chapters 5 and 6, since each goal addresses problematic aspects of those earlier works that constrained their development.

7.3 Preparation

Of course, to explore the nature of contact between human agents as it may occur through code mediated interaction, I needed to use technologies that would allow multiple human agents to each imbue a shard of their agency into the work, and so automate their agency. While my previous coding projects had used frameworks narrowly-oriented towards music application development, in the present case I required an environment that could represent different human agents more fully, and provide a platform for contact between them. I was drawn to the characterization of video games as social spaces (see Section 2.6.3), where a player might enter a game world as an avatar and then form a social identity. I have already discussed the concept of video game avatars, where I noted the capability of video game designers to form sites where human agents can engage in virtual interactions that change social relations in the real world (see Section 2.6.5). If human agents can meaningfully affect one another, and be affected by one another, through contact mediated by a virtual world, then a video game could provide a fitting model for my current artistic challenge.

7.4 Inspirations

In my characterization of locative media in Section 2.7.11, I describe a subcategory of the form, the “historical” variant of the “Placed sound” category. Works with such properties, like “murmur” (see Section 2.8.11), can reveal unknown or unappreciated histories of a place. Janet Cardiff’s “audio walks” (Section 2.7.10) also explore the idea of revealing a place’s meaning by adding a layer of information, which in her early work was delivered as an audio recording played by a Walkman. In both of these approaches, I felt there could be an opportunity to facilitate the development of empathetic bonds among cohabitants of a space, even if a participant was not physically present, nor personally known to the listener. This bond could arise, I reasoned, where a human agent interacts with a digital representation, or trace, of another agent. Locative media provides a platform in which augmented reality intervenes between user and space, populating it with artifacts that I thought might represent the experiences of other agents, and so form points of contact between

people.

As in “The Rough Mile” (see Section 2.8.11), where interaction with the app encourages its interactors to connect with a friend through experiencing a narrative, I felt that personal narratives could provide an object promoting empathetic engagement. That work combines aspects of audio walks and locative media with the aim of magnifying social bonds between pairs of people in the real world. Similarly, the makers of “murmur” saw the interactor’s encounter with unknown histories in a known space as a invitation to reconcile personal experience of a space with those of others, a process through which new meanings related to the space and its inhabitants may be formed. In each case, meaning comes from the augmentation of space. What might such a space look like, if it were to contain a multitude of distinct stories?

An answer appeared to me after I encountered the video game Pokémon Go, which, as described in Section 2.8.11, which was a groundbreaking, popular video game drawing upon techniques of locative media and augmented reality. When I played Pokémon Go for the first time, I was, as I imagine many people were, amazed by the introduction of the animated Pokémon into the world with which I was already so familiar. The initial surprise of sitting in a restaurant, a location I had frequently visited before and where I first installed the game on my iPhone, and the sudden appearance of a little monster that walked on the table in front of me, made me feel joy. The game had added a vivid layer to reality, the augmentation, that I had not imagined before. Chasing Pokémon became the charting of a new, previously-unseen world that existed within the real one—it was like discovering a treasure hidden in plain sight. And, like Schafer’s soundscapes, where the environment and the audio are inextricably associated, there appeared in Pokémon Go to be a link between Pokémon and the world around me. The first Pokémon I saw was not hovering in space, but rather walking on the table, which meant the augmented world obeyed the physicality of the authentic landscape.

As I considered the formation of and interplay between agencies occurring in my own algorithmic music, my experience with the Pokémon augmentations triggered a line of thought: the little monsters could seemingly act in the real world, be pursued and caught, while representing something of value. What could be of value in the environment, that an app could reveal? I thought of CPP, and how the museum wanted to use interactive music to help reveal the history of the Pictograph. I wondered what

other histories might exist around me.

My answer, simply put, is that everyone's history is around us. Like the Pokémon, human experiences are hidden throughout the world that humans share, unavailable unless a technology could be devised to make them visible. I recalled the many memories I had formed on the University of Regina campus since my childhood visiting my father's lab—they were meaningful to me, and, in a sense, they formed me. What if my app could evoke stories of people's real experiences, and use music to reinforce the connection between app user and memory contributor? It would take the core mechanic of Pokémon Go, but instead of catching Pokémon and training them to battle, the app would reveal histories that, once found and experienced, deepened one's appreciation for both the place they occupied and the people with whom they shared it.

"Shards of Memory," then, would be an archive of personal histories, encountered as an augmentation of real space, scored with music that reflected the interplay of app user, memory contributor, and myself as a composer and coder of its dynamic score. The meaning of the experience would derive not only from language, which Foucault has described as a filter that distances the listener from the reality that includes the speaker, but from music. I reasoned that if the app user has to expend effort to discover the memories, and could interact with them to contribute agency to their telling, this could form the basis of a social relationship modelled on those formed between video game players and their avatars (as introduced in Section 2.6.5). "Shards of Memory" would thus form a social structure. The user, whose active agency guides the unfolding audio-visual composition, would form the animating force. The memory contributor, whose agency is frozen like the inert being-in-itself described by Sartre, would ground the narrative history. The designer of the work (myself acting in blended roles of composer and programmer), would facilitate contact between the other agents, through coding an environment where they could encounter one another, through a joining of narrative and algorithmic score, that would deepen the meaning of that contact, a function that recalls the interpretive and strengthening role of a film score composer, a role with which I am well-acquainted.

7.5 Designing “Shards of Memory”

“Shards of Memory” is a work of locative media, using the placement of digital objects in an audiovisual augmentation of the campus of the University of Regina. The work is realized in an app running in Apple’s iPhone operating system iOS, and makes use of iPhones’ GPS sensors and the forward camera feed to allow the user to experience a superimposed representation of the campus, where I place virtual objects. I call the campus space the “real world” and the augmented space the “game world,” in reference to its construction in software typically associated with video games. The real and game worlds are mapped onto each other via analysis of the user’s GPS position, so that when the user walks a certain number of meters in a particular direction, the game world also shifts the origin of the point-of-view of the user proportionally, meaning that the relative positions of objects in the game world appear to exist at fixed GPS coordinates. The effect of this dependence on game world position on real world location is that they are spatially synced, as seen in Pokémon Go.

7.5.1 Technologies

The following technologies were used to develop “Shards of Memory” app:

- Game engine: Unity and the ARFoundation framework.
- Unity asset: ARLocation, created by Daniel Fortes. This asset aids in synchronizing user GPS position with a Unity-based game world.
- Unity asset: MIDI Player Toolkit, created by Thierry Bachmann. This asset enables script-driven midi playback through SoundFonts¹.
- Development platform: MacBook Pro (2013).
- Test platform: Apple iPhone X.

1. SoundFonts are sample-based midi-controlled instruments that may be played by a player such as the Midi-Player Toolkit asset that I imposed into my Unity project. SoundFonts typically contain recordings of single notes played by a musician on the instrument being sampled, and these recordings can be played back by being triggered by midi messages to create the illusion that the sampled instrument is being played.

The choice to use Unity as the engine of the app reflected my interest in the social aspect of the artwork. The frameworks that I had been using to this point (e.g. MaxMSP, pure data, Ableton, Arduino) offered wonderful affordances for manipulating audio and midi, but I felt that in order to create an experience that blended agencies and fused the real world with the world of memories that I envisioned, I would benefit from adopting technologies used to create video games. Unity offered a framework for building virtual environments and representing human social interactions, whereas the earlier tools, though powerful, were, I felt, too narrowly focused on audio manipulation. Given that I wanted my app to be the foundation of a social interplay, and understanding that Unity had been used to create many games through which people could encounter one another and interesting virtual characters², I decided to develop “Shards of Memory” in Unity.

I chose iOS as the platform for the app because of its compatibility with Unity, and because the iPhone granted easy access to users’ GPS location data. There were also additional audio and midi capabilities afforded by plugins available in the Unity store, which would prove necessary to developing the interactive score.

7.5.2 Setting and Gameplay

The world of Shards of Memory consists of a sequence of several phases: the Prelude phase, the Dream phase, and interlude called the “Tunnel,” and the Memory phase. The gameworld in each of these phases is populated with audio and visual objects that provide the experience of the app with purpose and depth. I will now describe the key objects that a user will encounter when interacting with the app, followed by a description of the stages of the experience that the user will encounter.

7.5.3 Audio and Visual Content

The content of the work consists of visual objects, and audio objects, and processes. All of these can be revealed when the user runs the app while physically located

2. In my own gaming experience, I have played the RPG “Pillars of Eternity,” and am a fan of Gareth Coker’s sublime music for “Ori and the Blind Forest”, and both projects were realized in Unity

within the Academic Green, a large, open, grassy space in the centre of the University of Regina campus forming the bounds of the content of the artwork. The visual objects are superimposed on the real world landscape at particular GPS coordinates. These objects, when approached, trigger associated audio elements. The following sections describe these virtual auditory and visual objects:

The Disembodied Voice. The Disembodied Voice acts as the user's guide and storyteller. The source of the Disembodied Voice is not identified in this version of "Shards of Memory," but it appears outwardly benevolent, though in its cadence there is a hint of danger. The Disembodied Voice was performed by Max Ferguson.

Blue Morpho Butterflies. The iridescent, crystalline Butterflies that appear in all phases of "Shards of Memory" were imported from a Unity Asset containing animated representations of blue morpho butterflies. I applied a material to their skin that rendered them eerily translucent. The Butterflies are enigmatic and I left their reason for being in the environment largely unexplained, allowing the user to come up with their own interpretation of the Butterflies' purpose.

Dreams. The Dreams hang in space in the Dream phase, and are used to help the user to interpret the melody that they will take with them to the Memory phase. The Dreams are static images depicting imagery intended to evoke a particular emotional reaction. I drew all of the Dreams' artwork on my iPad using the Affinity Photo app and an Apple Pencil in the months leading up to the completion of the app. The digital drawings were saved as .PNG image files, imported into the Unity project, and converted into sprites that could be mapped onto the surface of two-dimensional planes instantiated in the gameworld as virtual objects.

Melodic Avatar: A melody picked by the user as a reflection of their current emotional state. The Melodic Avatar is selected by the user in the Dream World and travels with her through the Tunnel Transition. It serves as a musical mediator between the user and the musical harmonies that surround each Memory.

Memory: A memory consists of visual and auditory elements. At the base is a Butterfly, rooted to a specific location in the vicinity of the Academic Green (specified by GPS coordinates derived from the memory contributors' description of their memory). A blue orb hovers over the butterfly, serving as a beacon that helps users to more easily find the butterflies. The memory is also associated with auditory effects the user may observe as they approach it. The closer a user gets to a memory,

the more the emotional qualities that the memory contributor associates with their memory begin to influence the ambient score heard by the user. A negative memory may be associated with minor keys or dissonance, and a positive memory may be associated with major keys and more consonant harmonies.

I will return to the structure of the sound world of “Shards of Memory,” which I will describe in more detail in the next section. Before that, however, I will provide a description of the phases through which the user will progress as they experience the app.

The Prelude phase

The user finds herself in a sparse simulation of a bedroom, consisting of a bed that appears beneath her, superimposed on the real-world ground. As with all of the objects the players encounter in the augmented reality layer of the experience, the objects are layered over the real world using computational processes afforded by the AR Foundation framework I integrated into the Unity project. Any AR object can be seen by the user, by pointing the forward-facing iPhone camera towards the real-world location where the virtual object is situated, and the virtual object will appear on the iPhone screen, superimposed on whatever real-world surface it happens to be placed on. In the Prelude phase, the user finds themselves hovering over the 3d-rendering of the bed, which can be seen as soon as they point the phone downwards and look at the screen.

The user may soon notice that a series of orbs appear floating around their head, and, if they look closely, they will notice an animated model of an iridescent blue morpho butterfly sitting quietly on the bed, its wings gently flexing. The significance of the blue morpho is not yet apparent, but it foreshadows elements of the environment that appear in later phases. The user soon hears some linear score, which I composed in the form of a lullaby, to suggest that, in the gameworld at least, the user is falling asleep. An ethereal Disembodied Voice soon begins to speak (the voice actor was Max Ferguson), welcoming the user into the land, and implying that she is already familiar with the strange dreamworld that she is about to enter. Once the Disembodied Voice has invited the user back into the dreamworld, the user may look down and see the bed fall down from beneath her, into the void below. The background, which has to this point been an inky blackness, begins to glow blue, and in the far distance a

column of rising clouds may be seen, surrounding the user. The Dream phase has begun.

The Dream Phase

As the Dream phase begins, the Disembodied Voice speaks again, welcoming her back to her strange dreamworld, and it proceeds to provide some context for the game world³. The Disembodied Voice reminds the user that she is going to experience the dreams of others, and after the voice concludes, the user sees several virtual objects appear nearby: shimmering, crystalline butterflies that walk and flit around the user, and a series of circular images which, the user is told, represent the dreams of people sleeping far below on the Earth. The voice presents users with a choice: they can remain in the Dream Phase and look upon the drawings, or they can “absorb” one of them, and then travel to the real world (phase 2) to encounter memories, which, unlike dreams, have a tangibility that the voice thinks the user wishes to experience.

Approaching an object is accomplished simply by walking, since the app is aware of the user’s changing GPS coordinates (updated every five seconds through checking the iPhone’s integrated GPS system) and each object in the game world is also placed at a fixed GPS location, which allows objects in the gameworld world to appear placed at particular locations in the real world. When the user enters the Dream world, she hears a looping musical background drone of choral effects I assembled from sample libraries I use as a composer. Walking towards a dream adds complexity to the music, for the dreams are associated with images that I have rated on their emotional positivity and intensity, and those ratings are used to modify the melodic, harmonic, and tempo features of the ambient score as the user approaches these Dreams. It is as if the area around a Dream causes a disturbance in the score, and the choral loop gives way to melodies and chordal harmonies that reflect the emotional valence and intensity of the image. The closer a user approaches a dream, the more strongly the harmonies associated with that dream will assert themselves in the music unfolding in the immediate geographic area. When a user is standing within one meter of a Dream, a melody will play. The melody is generated algorithmically based upon the

3. The Disembodied Voice assumes that the user’s character in the gameworld is the sleeper, not the user. Hence its dialogue helps a user to understand what is happening around them, even if they have never interacted with the app before.

emotional intensity and valence of the image. The user's task is to find a Dream for which the music, particularly the melody, resonates with them in some way. The Disembodied Voice informs the user that once they have lingered near a dream for long enough, they will "absorb" the music of the dream, which means that they will have chosen its melody to represent their feelings at that point in time. When they linger near a Dream for ten seconds, the Disembodied Voice tells them that they may look down to the Earth below, and they will then travel there to experience a trace of actual experiences (as opposed to the fictional Dreams of the Dream phase). The user looks down for five more seconds while standing next to a dream, and the process of transition is triggered.

The Tunnel

To reach the authentic world, the user needs to depart the Dream world, and they do this through a Tunnel formed of the column of clouds that bounds the horizon of the Dream world. As the user gazes down, the tunnel's initially gentle upwards flow becomes a torrent, emitting a screaming rush of sound, and the Dreams and Butterflies fall tumbling towards the ground as the tunnel roars. The end of the tunnel appears approximately eight seconds after the transition begins. The rush of air slows and the clouds dissipate. For the first time, the user can see a live camera feed displayed on the iPhone's forward-facing camera, with new virtual objects superimposed upon the physical landscape of the Academic Green. The user has entered the Memory phase.

The Memory Phase

Upon entering this phase, the user may notice that the melody associated with the Dream they stood beside before initiating the Tunnel transition remains with her. This melody has become what I term a "Melodic Avatar," and its looping theme forms the melodic content of this phase. As the user looks around, she may spy butterflies similar to those seen in in the Prelude and Dream phases. However, they are now joined by sparkling orbs that orbit above them, and approaching them now causes a different, more elaborate effect. Upon approaching a butterfly, the harmonies heard by the user begin to change. The background musical drone begins to feature

orchestrated chord sequences, and the harmonies and their orchestrations become more intense as the user approaches a Butterfly. The Melodic Avatar also begins to change. While the pitch contours of the Melodic Avatar are retained (preserving a degree of its recognizability as a specific melody), the pitches to shift, conforming to the scale of the harmonies that are associated with the Butterfly. When the user moves within one meter of the Butterfly, its significance is revealed. The Disembodied Voice begins to tell a true story, a Memory which happened to a Memory Contributor who was located in that very space when something happened that they chose to contribute to this app.

Each Butterfly is associated with one Memory, and the core objective of the app is now clear: the user is invited to encounter Memories of others with whom they share the campus space. The memory recitation offers a narrative account of the events surrounding the Memory, but this empathetic connection is deepened by the behaviour of the music as the user approaches the memory. In the Memory phase, they carry a trace of their own personality, represented by the Melodic Avatar, which was chosen during the Dream phase based upon its resonance with the user. In the conforming of the Melodic Avatar's form to the harmonies that surround the Memory, and the conforming of those harmonies to the emotional valence and intensity of the Memory Contributor's experience of the Memory during its formation, the algorithmic music of Shards of Memory serves as an intermediate point-of-contact between the user, the memory contributor, and myself, acting as programmer, composer, and facilitator of a new empathetic understanding on the part of the user for the Memory Contributor.

Once a user has heard a complete recitation of the Memory's content, the Butterfly flies back up to the sky, perhaps to join the Dreams. After encountering all Memories in the space, the app is concluded.

The dance of human agencies, richer than in any of my previous works and sustained primarily through the effect of an algorithmic, locative score, is a fulfillment of some of my key goals as a doctoral student. As such, I will elaborate on this system of agencies in the following section.

7.5.4 Agents And Agency

My aim for users of Shards of Memory is not to create conflict or to test skill, as often occurs in video games, nor is it to form a hierarchy of unequal participants in musicking as Christopher Small observes in Western classical concerts. My approach is to define agentic roles in such a way that they encounter one another in a spirit of mutual respect and understanding. I intend to raise empathic connections by allowing human agents to encounter one another. This idea has been introduced in Section 2.6.3, where I discuss Simon Waters' conception of the "Performance Ecosystem." Like such a construct, Shards of Memory blurs the musical roles of the agents, allowing them to each influence the form of the music, and indeed music here forms the tissue through which they are connected.

The four primary forms of agency that I identify in this artwork are:

- User, who triggers events and is represented musically through the Melodic Avatar;
- Memory Contributor, who provides Memories and so defines localized musical conditions during the Memory phase;
- Composer/programmer, a facilitator of agentic interplay by defining the form and function of the app; and
- The code itself, which can act autonomously to a degree, which I understand to be an autonomous source of agency.

The following sections describe each form of agency, while introducing my approach to generating particular aspects of the overall score.

The User (and Melodic Content Generation)

While developing Shards of Memory, I felt that I needed a digital representation of the user that recalled the video game avatars that I grew to love, such as the nameless hero of the Ultima RPG series by Richard Garriott (referred to as "the Avatar" by non-player characters), to more recent avatars like Geralt, the monster-hunter of the Witcher RPG series. I also noted how massively-multiplayer online roleplaying games,

such as World of Warcraft, form a ground where interactions between players and non-player characters can blur the distinction between real and artificial representations of humans (as previously discussed in Section 2.6.2).

The Melodic Avatar has already been introduced and described, but further elaboration of its creation and behaviour here will help to affirm its function within the app's interplay of agencies. It takes the form of a short melody that is regenerated in the Dream phase every time a user approaches a Dream. It is formed as a monophonic sequence of midi notes (each of which has a specific pitch, velocity, and length specification⁴) that can be transposed or otherwise altered in real time by the script that governs the midi-based music generation engine.

The music-generation engine is modelled upon my own composing methods. To model my personal composing process for integration into the app, I created a flowchart representing a simplified abstraction of my process for scoring a film. Then, I created a second flowchart depicting the flow of information in a generative music engine. Portions of this second flowchart were used as a template for coding the generative music system. In particular, a section of my flowchart outlining three distinct processes I tend to use for composing melodies became the basis for the generation of melodic avatars. These procedures are outlined in Appendix C. They serve as a tangible point of contact between my action as an agent composing a score for a film, and my agency as projected through Shards of Memory's music generation engine. In the correlations between my film composing process and the blueprint for designing the music generation engine, there is a correspondence between my agency as applied in producing scores I have composed for film, and my agency as it acts within the musical-making algorithms I coded for the Shards of Memory app. This is a key connection that I will return to in the reflection.

While melodies are generated in proximity to a dream that has a certain profile for emotional valence and intensity, each generated melody must be capable of conforming to any possible combination of valence and intensity. Thus, all melodies are initially composed in the key of C-Major, but when they follow the user into zones

4. Pitch is represented as an integer number from 0 to 127, numbered in semitone increments with middle C assigned the number 60. Velocity here corresponds to the loudness of a pitch, on a scale of 0 to 127 with 1 being the softest and 127 the loudest. A velocity of 0 indicates that a note should stop entirely. The length is an integer representing the number of eighth notes duration of the sustained note. The midi playback engine I designed has a resolution of one eighth note, and since tempo is variable, the timing of an individual eighth note is not a constant.

specifying other harmonies, they are conformed to those harmonies by a function that filters notes. This function is a more complex implementation of the pitch-conforming system I used in Moonloops, expanded with a larger set of possible chords. A single melody generated by this system can be conformed to any combination of intensity or valence rating.

Generation of Melodic Avatars involves my agency in two processes: production and evaluation. Rather than write a complex algorithm to generate a single, subjectively-well-formed melody, I envisioned a design that modelled not just my compositional processes, but also my musical preferences. To this end, I coded three simple systems for generating melodies, and, every time a new melody was required, the system would create ten variants. The three systems were based upon my own personal means of forming melodies as outlined in the Film Scoring flowcharts⁵. These ten variants were then fed into another function which rated each prospective melody on how well it fulfilled three criteria:

- The degree of conformance of the rhythm of the melody to a simple march rhythm I commonly use in my music;
- The appropriateness of the density of note events in relation to the affective conditions near the dream where the memory formed;
- The tendency of the melody to feature small changes in melodic contour rather than large leaps, reflecting a common feature of many of my melodies.

The melody that scores the highest overall rating based on these three criteria becomes the melody played when the user approaches its associated dream. Thus, I maintain that my agency may be considered active both in the composition of potential melodic avatars and in their evaluation, and so my agency as a composer is embedded in some degree into the behaviour of the melody-generation code.

The instrumental sound through which the Melodic Avatar plays is chosen randomly from a list of available SoundFonts that I imported as assets into the Unity project. I needed to balance embedding my agency in the music generation system with allowing enough control by the user to give them a sense of investment in their

5. See flowcharts 4.1.1, 4.1.2, and 4.1.3 for descriptions of how the generative systems work and how they are derived from my film composition behaviours.

Melodic Avatar and to feel that they had some control over in-app events. Randomizing the instrument was part of my attempt to strike such a balance—while I have clear ideas regarding instrument choice in my composing practice, I felt that specifying instruments would diminish the opportunity for users to find music that suited their own taste. As such, I randomized both the instrumental sounds of the avatar and the system specifying pitches in the Melodic Avatar generating function. The result is that I cannot fully predict exactly what a possible Melodic Avatar might sound like, while the user may still find melodic structures more instances of a melody that they feel reflects their mood.

Memory Contributor (and Harmony Generation)

The Memory Contributors provide text-based Memories that reflect particular, memorable experiences that have some significance to them. These are gathered as “Memory Objects,” which include a few sentences describing a memory and some additional data, all of which are defined as follows:

- *Memory Text.* A textual account in English of something that happened within the bounds of the artwork, limited to three sentences. Contributors, who remain anonymous as they may be contributing personal information, were encouraged to be honest, to not identify others, and to avoid vulgarity. Other than that, they were free to contribute any memory that they wished to freely share, so long as they were present when the memory formed and involved in the events to which the Memory Text refers in some manner significant to them.
- *GPS coordinates.* A precise location where the events described took place, that could be represented via GPS coordinates. Where the contributor was unable to provide precise coordinates, a textual description of the location was used to infer the GPS coordinates.
- *Emotional Valence Rating.* A rating on a seven point scale, where a rating of “1” was associated with extreme displeasure, and “7” with extreme pleasure. A neutral rating of “4” could indicate ambivalence or indifference. Crucially, this rating is based upon the emotional state of the Memory Contributor in

response to the event itself, not those of others involved nor of their current feelings regarding the event.

- *Emotional Intensity Rating.* A rating of the intensity of the emotion the Memory Contributor experienced while participating in the actions described in the Memory Text. A score of “1” could indicate numbness, which increased up to the rating of “7” which might denote extremely intense emotional experiences.

The following is a representation of a Memory Object, used in the app and provided by an anonymous contributor. All content was provided by the memory contributor, except for latitude and longitude, which I determined in this case by consulting a mapping app on my iPad and identifying the CPS coordinates of the brick kilns on southern boundary of the Academic Green.

- Positivity Rating: 6
- Intensity Rating: 5
- Location: The brick kilns, out in the back of the art department in Riddell
- Latitude: 50.4174, Longitude: -104.5907
- What Happened (quoted from the Memory Contributor): There were many times where I would close my eyes in that space. It was fall, and the smell of earth and crackling twigs, wood smoke, and oxidizing metal floated on the breeze. It took me back to my childhood camp days, watching that fire. When I closed my eyes, for one moment, all was good, peaceful and simple. No human complications, no pain, no fear of the future.

A brief comment on the representation of emotion in these is warranted here. The description of the relationship between human emotion and music has been an ongoing concern for psychologists since at least the 1960s (Eerola 2018, 540). My goal with Shards of Memory is not to add to this research, but rather leverage some of its findings to enhance empathetic connections that form between agents who act within its structure. There are two primary studies that inform the creation of the memory representation system I describe here.

The first is James A. Russell’s “Circumplex Model of Affect” (Russell 1980). In this model, emotions are identified along two axes: pleasure-displeasure, and arousal (1167). Russell argued that these axes formed a perimeter where opposing emotions (e.g. sleepy versus alarmed, or sad versus happy) were located 180 degrees apart. This model is relatively simple, in that it can map most emotions easily through numeric coordinates of arousal (which I termed intensity when working with Memory Contributors), and pleasure/displeasure (which I simply termed “positivity” or “valence”).

The second system is Ekman and Cordaro’s theory of basic emotions (Ekman and Cordaro 2011). They list emotions that have been empirically studied and found to be both discrete, in that they can be empirically distinguished from one another, and that they are found to have evolved as adaptations to our surroundings (364). This latter criterion points to these emotions as being universal, shared between humans of all cultures. This is not to say that the experience of these emotions is uniform, nor that every person has the capacity to experience them. However, in that these emotions are essential and close to universal, I feel that their identification in the metadata for contributed memories helps to clarify any ambiguity in the valence and arousal ratings.

Taken together, the ratings of emotional valence and intensity provide a simple means for contributors to accurately characterize their experience of their memory, while helping me, as the designer of an interactive, algorithmic score, to evoke the emotions associated with contributed Memories (the content of which I typically had no knowledge of when designing the generative system).

Memories and Dreams appearing in the game world affect the properties of the ongoing score when the user moves to within three meters of that object’s GPS coordinates. The particular chords that are evoked near the object are generated at runtime. Each Memory’s harmonic pattern is constructed by a function that takes as input the valence and intensity ratings specified by myself as the Dream contributor (during the Dream phase) or the memory contributor (during the Memory phase). The chord patterns are stored in a Scriptable Object which represents chords as a sequence variables of two values each. The first value is an integer representing the chord root (with the note C set to zero, $C\sharp$ set to one, and so on in ascending order of pitch), and the second represents a chord type (where a major chord is arbitrarily

defined as the integer “1,” minor as “2,” augmented as “3”, and diminished as “4”).

The harmonies and the melodies are played as a sequence of midi events. Regular timing was essential, to prevent jittering of the midi notes’ timings. I used the Unity Asset “Koreographer” to control timing. Koreographer allows audio file playback to trigger Unity script events, at timings tied to particular samples in the audio file. To create a continuous roll of midi events, I recorded a single bar of silence at 135 beats per minute and created eight events tied to samples corresponding to the samples at eighth-note intervals on the timeline of the audio clip. I set it to play back, and it became a metronome that triggered a event every eighth note. To produce tempo changes, I simply bent the pitch of that audio file upwards (to raise tempo) or downwards (to lower it), coding a small function that took a target tempo value as a parameter and then over a specified number of seconds bent the pitch of the audio file until event triggering aligned with the desired tempo. Since, in the vicinity of a Dream or Memory, musical content is entirely realized through midi-controlled SoundFont playback, tempo changes could occur smoothly and the midi timing was scaled appropriately. This feature was useful given that tempo is an important influence on affective responses to music.

Harmonies were realized through midi-controlled playback of custom SoundFonts, the audio content of which I recorded in my composition studio. I recorded two sets of chords, one for the Dream phase and the other for the Memory phase. In each set there were multiple instrumentations. For example, I separately recorded piano chords, low string chords, high string chords, and tonal percussion for the Memory phase, in several different keys and chord types. Because the custom SoundFonts had renditions of every possible chord types, root chord, and orchestration available on a single midi channel, combinations of these chords could be blended and sequenced to create the effect of a smooth change in texture complexity. On the edge of a Memory’s zone of influence, for example, only the low piano chords might sound in a minor key, but as the user advances closer to the Memory, other orchestral elements such as the strings might be mixed in. The orchestrations were thus layered, and reflected my personal cinematic compositional style. A single SoundFont patch could permit the app to trigger a harmonic progression comprising any combination of the available chord types, in any of the orchestrations I recorded, for any chord root of the chromatic scale, and with any of the four chord types I selected. This system

provided a great deal of flexibility in generating harmonies, which was helpful given the goal of generating music that effectively supported the emotions of contributed Memories.

Composer/programmer

The third source of agency that is active in the functioning of Shards of Memory is my own, which is projected principally through acts of music composition and programming. If my goal in creating Shards of Memory was to find a way to share agency over music in a manner that promotes empathetic connections, I must acknowledge that my own agency, in devising the interface and audiovisual content, is still quite powerful. I designed the functioning of the melody and harmony generation systems, even though I allowed entry points through which a user or memory contributor could affect the unfolding of the music. I composed the orchestrations of the supporting chords, even though the timing and mix of these elements was dependent upon the behaviour of the user and the nature of Memory Contributors' narratives. On a more elemental level, the concept underlying Shards of Memory—that a locative, generated score can help to form empathetic bridges between agents outside of its designer, who may empathize with each other through engagement with the app—is my own invention, and this purpose cannot be modified by users or other agents during runtime. Because I did not relinquish control over the unfolding of Shards of Memory more dramatically, my own agency remained privileged in some significant ways. The behaviour of the app is constrained by my taste in music, and my values. For example, since I personally have little experience in composing atonal music, the melody-generation system is not designed to produce them, and a user who has an affinity for such music might not find the Melodic Avatars as satisfying. Such limitations are the imprint of my compositional agency realized through my practice of coding. I consider this aspect of the app a strength as well as a weakness. It is a weakness if it prevents a user or Memory Contributor from feeling satisfied that the generative score reflects their personality or state of mind. It is a strength, however, in that my own agency, informed by a lifetime of composing scores that often support others' narratives, can be directed towards composition of themes that are effective and compelling to listeners. My own empathy is engaged when I score a

film, and I feel the same empathy helped me to design the generative engine, the custom SoundFont-encoded orchestrations for the harmonic material, and the conditions under which these score elements would be presented. That said, my agency was absolutely required, since if the onus for specifying musical events was wholly placed on users and memory contributors, especially those with minimal musical education or instrumental performance experience, it is hard to imagine how such individuals would create music that resonated with them. So, my choice to take control of certain musical variables and functions, but not others, helps the score to function as intended, and effectively. In future work, I will explore other paradigms by which to balance the power of agents taking on different roles in algorithmic-music-based environments, but at this point I feel quite comfortable claiming some compositional control, recognizing its value where musical scores are automatically generated.

7.6 Interplays of Agencies in the Algorithmic Music

In some dynamic video game scores (as outlined in Section 2.7.8), music can be subject to conscious or unconscious control by the player. The music is nonlinear, even though it may appear linear to the listener on a single play-through. The music and the speech renditions of the Memory Text form the “Locative Score,” the form of which is governed by three human agencies:

7.6.1 Divisions of Control

I determined that I needed to delineate that which belonged to the user from that which belonged to the Memory Contributor, and myself. To this end, I decided to create a score in which melody fell under the purview of the user via their chosen Melodic Avatar. Harmony, tempo, dynamics, and the spoken memory narrations were sourced from the Memory Metadata provided by the Memory Contributors. My own role was to fill in the remaining components of the score, such as its arrangement, instrumentation, and the event systems that triggered changes in the score during its generation, so that the emotional connotations of the overall music supported the affective character associated with nearby memories, and that the Melodic Avatar would retain its identity while duetting with thematic conditions that appears when one nears a memory.

This interplay of agencies required me to design an adaptive music system that both served the two agent classes external to me, while still allowing me to join in the musicking in a meaningful way. In the next section, I describe how I leveraged my experience as a film composer to animate my agency in the Algorithmic Score.

7.6.2 Conceiving the Algorithmic Score

There are a number of lenses through which I might have framed this multi-agency, music-mediated experience. It could have been positioned as a group improvisation, where users, memory contributors, and myself make music unbounded by the more restrictive requirements of the notated score. However, the distinction between roles in *Shards of Memory* suggests that this lens may not be the best fit, as there is an inbuilt imbalance between users', memory contributors' and my own means of controlling the unfolding of the experience, rather than a sense of equality. Another lens might be the concert, where the user becomes the audience, memory contributors resemble musicians, and I would be the composer and conductor, creating raw music which the memory contributors shape based upon the emotional connotations of their memories. Of course, remaking the concert hall in an augmented reality artwork may merely reintroduce the classical concert hierarchy, a position, I have argued, where my agency is set apart from the others'. As such, I rejected this model as well. I needed a model of interrelating agencies with which I had some experience, and which could admit the action of all participants in the making of the experience of the app.

I realized that in approaching this algorithmic score, my greatest asset, so to speak, was my experience as a film composer. As has already been stated, the design of *Shards of Memory* was modelled on video games—avatars, adaptive music, locative media, and the coding environment are all influenced by contemporary video game design. I considered how I would score a video game (at the time, I had not yet had that opportunity), and looked to my favourite scores. In particular, I looked to Jeremy Soule's score for "Skyrim," (Bethesda Game Studios 2015) an expansive, richly cinematic, and affecting score which has become a personal favourite game score. It was not strictly a generative score, but rather was scored for a symphony orchestra in the manner of some of my favourite scores from my childhood (in particular those of John Williams). I wondered if I might be able to model my agency

by drawing on my particular approach to cinematic scoring, as I consider this ability to be one of my compositional strengths. It would also form a new point of contact between my dual identities as a soundtrack composer and as a coder of algorithmic music, since none of my previous doctoral projects had built upon that aspect of my compositional practice. I decided that *Shards of Memory* should feature a cinematic, interactive score.

On first consideration, a film score may seem quite contrary to the mobile, layered world of augmented reality—to me, the experience of film scores recalls a concert hall in its division of audience from the unfolding of events being projected. Yet beyond the context in which we experience film and television, there is another way of understanding their nature. I see watching a TV show or film like being in a hybrid space. We sit, sometimes with trancelike focus, physically present within the familiar surroundings of home or theatre, gazing into a glowing screen. When I am absorbed in a show, it fosters a certain detachment—I am both inside and outside of the fictions that the screen and speakers offer. I watch but do not participate, sometimes reaching a state of detachment that makes me porous to the projections of others. It can be like limbo, where time stands still because you are not really present in your living room. As a composer of soundtracks for film and television, I feel partly responsible for shepherding the viewer into that space between chair and screen. As a composer for such media, I seek to maintain the viewer's connection with the imagined world, linking events on the screen with particular emotions. If I am successful, the music helps to bridge the viewer and the view, so that the characters acting on the two-dimensional screen seem real, and worthy of our empathy. This goal is consistent with my aims in *Shards of Memory*, and I wondered if I could use the composition of a film score as a model for the generative music system in the app.

A film score is fixed to its film, and while I am experienced in composing for film, the current task necessitated that I identify a tangible link between my action as a composer and the behaviour of the algorithms that would realize the interactive score for *Shards of Memory*. While Soule's score to *Skyrim* is a popular, evocative video game soundtrack, it is also essentially a linear score, since many of the individual cues are presented as orchestral recordings. There is nothing inherently wrong with such an approach, but in the current context, where music animated by code helps to animate human empathy, I could not simply compose a score for each memory. I

wanted the music to dance with and between the agencies whose touch passes through it. What I needed was an agentic representation of myself in code, that would behave as I would were I scoring the real time unfolding of *Shards of Memory* myself. Such an agent, if it were a mirror of myself, would allow me to project myself into the score and musick with the other agents. In this act, I would fulfill the missing element of Q-Chords, Moonloops, and CPP—an embrace of contact between myself and the others, through music that would reflect our agencies. The algorithmic engine that generated the music of *Shards of Memory* would be based on an algorithmic depiction of my process for composing a film score.

I created two parallel charts, which I call “Paths of Agency”. The first, labeled “Film Score Path” represents the series of processes in which I engage when scoring a film. It is a general depiction, not specific to any single project. The second is called the “Generative Path.” It represents a translation of the processes and decisions indicated in the “Film Score Path,” to a form compatible with a hierarchy of objects in an application coded in the Unity framework. The algorithms governing the behaviour of the music of “*Shards of Memory*,” including the flow of control and the particular inputs and outputs that form points-of-contact between the algorithm’s functioning and the sensor inputs and audiovisual outputs supported by the iPhone, are indicated in the structure of the “Generative Path.” These charts are included in Appendix C.

My argument that my agency is active in the functioning of *Shards of Memory* is underlined by the links between the two series of flowcharts. Through designing an algorithmic music engine modelled on my methods for composing a film score, I am able to demonstrate that I am an agent in *Shards of Memory*, in the form of a generative music structure that acts as I would if I were to score the events that unfold in the execution of the app as a film. As such, my conclusions on the nature of agency revealed through the chronological sequence of programming projects I have developed during my doctoral program are informed by the experience of breaking of a shard of my compositional agency, and letting it animate the behaviour of the app, serving as a point of contact between myself and the other two forms of human agency that operate there.

7.7 A Reflection on Mingling Agencies

Insofar as the generative music engine for *Shards of Memory* is scoring a space, it is intended to fulfill similar goals to those of many of my film and tv scores: to take representations of human activity and make them matter more to other people. In a film score, it is the personality of the characters, and what happens to them as they act out their narrative, that I pay the most attention to when scoring a scene. In *Shards of Memory*, I have access to similar constructs—the memories reveal events which reflect their contributors’ personalities, in the metadata values they associate with them, and in their choice of what is significant enough to share, and what means less to them. In this sense, I am treating the memories like the characters and plot in a film, where my adaptive score serves to magnify the resonance of these events for the user. The users themselves are like an audience, but an audience who is absorbed into the telling of a story, to where they feel as if they are part of its telling. This absorption of the users is based upon the hybrid space that I described earlier, the limbo between reality and a compelling fiction, a state to which I always aim to help take the viewers of a film I score through the design of the music. In “*Shards of Memory*,” I wish to create a similar suspension, with the key difference that participation in the telling of the story is not a simulation, as in the viewing of a film that conveys a pre-sequenced series of narrative events that we cannot change. Rather, the app’s use of code-generated dynamic music that responds to the user’s behaviour and state of mind as embodied in the *Melodic Avatar* can actualize the user’s participation in the unfolding of the events.

Shards Of Memory is my fullest exploration yet of the mingling of agencies in the context of the interactive new media artwork. It is the product of my desire to musick as an agent among agents, where my contact with those others is mediated by program code and algorithmic music. I view this form of social contact between agents as an act of vulnerability, where the coder detaches a sliver of her spirit through imbuing into code the honest imprint of her being, revealed when the code executes and acts in the world as a reflection of the coder. This is a gift that the coder can give, a piece of her self, which is arguably a person’s greatest treasure.

To Descartes our awareness of our self is the beginning of knowledge: “I think, therefore I am.” I feel that programming involves sharing an element of our self. The

self lives in the action of the autonomous agent that derives from our self, that mimics our behaviour without our control, a puppet without strings. To automate ourself in code is to introduce into the world another spirit, where a spirit is something that acts within and through a being, animating it. That is the difference between tool and algorithm, a hammer and an app. A tool is inert, an algorithm is not. In *Shards Of Memory* I finally embrace the work in its essential nature. It is an algorithm, acting among others, a product of my control, but in the end not fully governed by it.

I have framed the central problem in my doctoral studies as having been related to what I call a crisis of agency. For me, this crisis manifested in the form of a reluctance to relinquish control over audiences' experience of music on the one hand, and in my fears of computer-borne compositional agency somehow devaluing what I perceived to be an authentic "human" music, on the other. In both cases, I fearfully anticipated the devaluation of my music, first from a rising babble of scattered human agencies whose interactions with my music would precipitate its descent into noise, and second from the falling sword of some superior algorithmic composing system that would exceed human composer' compositional capacity, making our musical expressions redundant and meaningless. However, these twin anxieties, I feel, formed at a time when I saw music as a transmission, and myself as being apart from others involved in the realization of a performance of my music. If a computer could compose, I thought, naturally it would become a composer. And, if the teeming multitude of non-musicians were granted programmatic crutches enabling them to initiate the composition of music effortlessly, then why were composers special? Why was I special?

In both of these anxieties, I was adopting an hierarchical perspective. I saw the role of the composer as isolated in a position of control over music, and the agency of such a composer as exerting their power through the notated scores they write. I saw rigid boundaries between composer and musician and audience, as well as between human and machine. The cyborg, obliterating the ingrained boundaries between human, animal, and machine, awakened in me the notion that embracing the machine might actually serve to magnify my capability, but I was still hoarding responsibility for music, cyborg or no.

When I encountered Deleuze's rhizome, (see Section 2.3.3), and encountered Small's

inclusive behaviour he called musicking, I realized that my tendency to compartmentalize needed to be challenged. I recognized that music could be understood as a hierarchical construction in its social and structural forms, and I believe that to be a strength. After all, music's notational grammars and syntax can serve as a means of facilitating the communicative functions of music. The abstract nature of music may necessitate forming a mutual understanding of such grammars or conventions, in order to appear unequivocally and interpretable as music. In my case, the decentralization modelled in the rhizome or in the musicking community did not inspire a new form of music, but rather revealed to me a new model of the social milieu in which humans, and machines, might practice music meaningfully. This milieu is found in the space where humans encounter one another's agencies, mediated by a digital computer. I will illustrate this using the example of my app, and how agencies may contact one another through interacting with it:

In *Shards of Memory*, the Melodic Avatar provides the thematic content of the music, a monophonic melody conceived in the form loosely based on the Wagnerian Leitmotif that remains in use in film scores today. The memories provide the emotional foundation of the score, where increasing proximity to a memory magnifies the degree to which its contributor's affective ratings influence the expressive content of the unfolding score. The voice of the Memory Narrator speaks words linked to its contributor's emotional associations with the event, and the score should increasingly project these emotions as a user nears the GPS Location associated with the event. The goal of the generative engine, then, is to bring these musical elements together in an interplay that reveals to the users the meaning of the memory, fostering empathy in the user for the memory contributor. I feel it accomplishes this task most effectively when the users see themselves in their Melodic Avatar, then follow its entry into the zone where the memory begins to take control of the harmonies surrounding it, and finally experience a reflection of the emotions associated with a memory by an anonymous contributor. In this sense, the score is acting as an adaptive analogue to film score, or a video game score that responds not to in-game events, but to the emotional contrasts between its users.

In this light, we can understand how agencies may mingle in the space of the work. The memory contributors' agencies appear frozen in memories they have

shared, which manifest as augmentations of campus. The users' agencies are active, constantly projecting into the space the users are exploring. Mingling occurs in the zones of influence of the memories. As users approach a memory, their proximity to the memory begins to change the quality of the harmonic texture, like the gravity well of a black hole bending light. The Melodic Avatars engage with the memories, morphing in a duetting response to the memories' transformation of the Harmonic Environment. A memory's contributor is likely not physically present, but their frozen agency transforms the unfolding music in a way that impacts upon the user, musically and verbally. Like music added to a scene in a film that in its unscored form projects ambiguous emotions, the music evoked by the interplay between the Melodic Avatar (a representation of the user's affect) and the Harmonic Texture (with expression aligned to memory content) draws the user into the telling of the memory, magnifying its emotional expression. The combination of familiarity with the space, a narrative of a true event that impacted someone strongly, and music that reinforces the emotions reported by the memory contributor, will ideally create an empathetic understanding on the part of the user directed towards the anonymous contributor.

This bond, between user and memory contributor via the proxy of music and narrative, is formed as a consequence of active engagement between the two individuals, and is mediated by their agencies. This mingling of agencies, where frozen agency is shared via a narrative document and associated metadata, however, is in some ways a recapitulation of the transmissive composer/audience model of sharing agencies I identified in the classical concert. After all, the memory contributor need not be present in the space to exert their frozen agency, and if they are not involved with the user either while the app is running or after it is shut down, they cannot be influenced by the mingling.

The users, however, are no mere concertgoers. They are represented in both layers of the app's domain—as GPS-World people learning about the depth of experience enriching the world about themselves, and as an avatar-mediated participant in a digitally-constructed augmentation of that world, revealing its richer meanings. As such, they have agency over their experience, over its unfolding and affect, in a manner that erases the sense of separation between user and contributor.

So far, I have identified two sources of agency sustained in the domain of the app.

Of course, there is a third, that of my own agency, which I would like to outline here. In Appendix C I share flowcharts that show parallel paths between my general approach to scoring a film, and the structure of a hypothetical algorithmic music engine that functions similarly. These charts trace the behaviour of my agency, and, in parallel, they reveal the common purpose of my action as a composer of notated scores and as a coder of compositions. Since in *Shards of Memory* my agency has been sealed within the code, to be carried into the layered domain of the Game- and physical Worlds, I would argue that I have now uncovered a resolution to the issues that I believe diminished the effect of my earlier doctoral projects: in “*Shards of Memory*,” my agency does not retreat to hide behind the code, nor does it attempt to suppress others to the degree that my earlier works did. Instead it acts as I would, with respect for those other agents it contacts, using the sharing of control as a means of sharing the self. Though the experience of agents in *Shards of Memory* was mediated by programmatic automation, for the first time it appeared that I was present in the functioning of the app, and that was a revelation.

Chapter 8

Conclusion: On Agency and Automation

The works presented in this dissertation trace a process through which I explore the unstable nature of human agency where acts of composition are transplanted from functionally-linear notational media to functionally nonlinear algorithmic media. This body of work was carried out using interdisciplinary tools, techniques, and perspectives that drew from the discipline of computer science and from the European tradition of classical composition, in both of which I have a history of creative engagement. The four chapters documenting my creative work reveal my personal struggle to locate myself meaningfully in relation to other people through the development of algorithmic music systems. My goal as an algorithmic composer, however, remained essentially the same as my goal as a film composer and songwriter: I wished to magnify empathy and social bonds between human agents through musical expression. In *We, The Artificial*, I explored this idea through the metaphor of the self-aware but unacknowledged machine, where the words of androids and robots I wrote as lyrics mirrored my own sense of displacement from others. In the chronological sequence of the four remaining primary doctoral works—*Q-Chords*, *Moonloops*, *Chief Paskwa's Pictograph*, and *Shards of Memory*—I tested the bounds of my agency and its expression through code, and explored the meaning of its contact with other human agents. Taken chronologically, these algorithmic works revealed approaches for creating satisfying musical experiences that balance the interplay between agents, a condition that dispels much of the anxiety I have experienced since first imagining futures where automated composition devalues composers, as opposed to offering them extensions of their practice and expression.

8.1 Agency in Code and Score

I want to return to an argument I made in Section 4.5.2. I argued that a computer transcends a composing tool such as a pencil, or an unambiguous extension of human action, such as a hammer, because it carries out processes. I explained that when I use a digital audio workstation to compose music such as *We, The Artificial*, I am availing myself of processes (e.g. use of templates, presents, or parameter automation) that introduce other agencies into my workflow. This occurs because automation requires action towards a goal, and goal-directed action implies an intent, which must come from the designer of the code, wherever the code itself is understood to lack an innate intent of its own. Throughout my doctoral program, I had been trying to find my composerly self in the code I was creating, and so form a site where agencies—my own and others’—could come into contact. This was a gradual progression that I recognize in the changing capabilities and limitations of the series of works described in this thesis. I came to frame this search as a search for an interplay of agencies, borne by code that I created to blend my self-expression with others’, in a way that would transcend the linear scores I had long created with Cubase. By so doing, I wished to nudge the general practice of music forward in some manner, although, initially, I could not articulate what that manner was. What I realized as I was working on *Shards of Memory*, however, was that an organized, goal-driven process originating from a human mind can contain an intent, and that intent is the motivating force of the agency that acts in accordance with its aims.

When David Cope says that “all composers are algorithmic composers” (Cope 2015, 405), he is right—music is process-based, involving an enactment of choices that may reflect an intent. Now some, like Cage and Xenakis, have chosen to deny the necessity of intent as a form of direct specification of musical events, yielding to chance or other external agencies the control of musical variables. Yet such musical works are still processes, and they have some intent that propels their form. Even Cage’s *4’33”* needed a score, though it is minimal and textual, otherwise how would one have known that it was happening? Cage’s intent that shapes its presentations makes *4’33”* a composition, and that intent cannot exist without his agency, applied through the performers’ and audiences’ observance of the conditions of its unfolding, animating the musicians’ silence that marks its performances.

If Cage’s composerly agency can act without the sounding of a single note, perhaps musical agency need not be restricted to the constraints of a score. Indeed, I consider my composerly self, and the expression of my agency, to also be active when I code. The processes in which I engage as a composer share with those evident in my software-based works a personal intent that I claim as my own. My agency is the vessel through which that intent acts outside of me and so is the source of animation that sets the works I make in motion. In this sense, composing with code and through notation may be indistinguishable, in that both set into motion processes that carry my intent towards an audience, whom I hope will be transformed in the expression of their own intent, through their own actions, such as clapping, reconsidering a belief, or engaging with a wave of emotion that they otherwise would not have experienced. Where I receive a positive response to a performance of one of my compositions, as when the user of *Shards of Memory* experiences a pleasant realization about the space she occupies, I feel that essentially the same mechanic is at play—I express my intent through a medium that allows the manipulation of musical events, collecting them into a process that projects my agency in accordance with that intent, and applying my agency to affect those who encounter it by altering the unfolding of their own behaviour.

8.2 An Agentic Mirror

Why, then, does the idea of a composer acting through code instead of through notation matter? My facility with computer technology has not changed my means expression or intent, nor my values. The answer, I believe, lies not in the form of creation, but in the form of the interactions that result. Music, a temporal art, is necessarily a process. A process can exist in a spectrum of determinism, from being utterly linear and invariable, to utterly chaotic and unpredictable. The position of algorithmic music (by which I mean its interactive or adaptive varieties facilitated by programming) forms a range that might differ from the range upon which musical performances deriving from Western Tonal scores may be found. Neither approach sits at a pole of this spectrum, representing total chaos or complete determinism. But the ranges, though they may overlap, are distinguishable. The notated score found in Western Tonal music typically (though not always) imposes a template on the

performer, admitting expression within a constraint system specified by the events that the composer dictates through the score. The agency of the composer conveys her intent by creating a frozen structure from which restricted variation, the performer's expression, may emerge. This is why the performance of a particular work can vary and yet still retain its identity. The work is a combination of a deterministic score and a variable performance. When composing through code, on the other hand, an additional site of variation may be uncovered. That site is formed by the processes that the code carries out, and it exists in a layer of execution of the work that is, by analogy with the notated score, situated between the events inscribed on score paper and the performer's expression. There is typically no distance between the performer as the score in a solo instrumental performance, for example—the performer reads the notes and acts in ways informed her own interpretive processes, carried by her musical agency, in completing her intended realization of the score. I believe code resembles a score, as the two media appear in my own creative work, since the code is fully specified before its execution. Upon execution, however, the play of agencies in the two forms diverge.

In the case of a tonal musical performance, the performers (and possibly a conductor) exert some agency over the sound produced. They read the score directly, and, within the constraints imposed by their performance tradition and the structural dictates of the score, they can contribute their own expression. The notated score may be seen as a document of the composer's intent, and the performer works directly from it. Actions that animate the score are an expression of the agency of the performer, conductor, perhaps the audience. But there is no action on the part of the composer. Alive or dead, if she is not present, she cannot act in the space. In such a context, the composer's agency is a frozen agency. It is fixed in her score, and travels into the future on that basis.

As I extended my coding practice to the medium of code, I realized that, while both the programs and the notated scores I was created were frozen. the processes that emerged from the action of my code were of a different character. When such code acts, it remains under the control of the coder, since the act of coding specifies the action of the code. Where in my notated or recorded compositions, I was creating an inert document that could only act under the force of another's agency (e.g. to strike the piano keys or, more humbly, click the play button on a music app), the

action of code did not need other agencies to animate it. Consider George Lewis' *Voyager* (discussed in Section 2.8.1). It could improvise with a performer, or it could improvise solo. Where did the animation come from in the soloing mode? The code itself. Likewise, in *Shards of Memory*, where did the dancing melodic avatars come from? Again, the code. What would the equivalent be in the notated score? Any animation of the score, which we would call "the performance," is a result of the action of the performer. Superficially, these two animating forces seem to be equivalent, but they are not. Why? Because the animating force for the notated score is the agency of the performer. The animating force of the behaviour of the program is the code, and the code is the product of the agency of the coder. In that distinction, which became apparent to me as I considered the difference between my projection of my agency in *We, The Artificial* as opposed to *Shards of Memory*, I see music's future. The composer acting through code is granted action beyond the document of the code, for the code specifies behaviour that can conform to her intent by acting as an extension of her agency, independently of other agents. The notated score is inert but for the action of other agents. This means that coding a composition grants the coder an additional space of intent-driven action, one that does not exist in the chain of agencies involved in a performance of a Western Tonal notated score.

Let us call the intervening layer of processes that emerges from the coded specifications of a programmer, and are active before other agencies intervene in the system, the "agentic mirror" of the composer. What might the agentic mirror provide to the composer? For one thing, it can act independently of the composer, but preserve her intent. For example, in *Shards of Memory*, the augmented landscape incorporates memories that are the nexus points of a heterogenous score, in which the complexity of the score increases where memories are nearer to the user. That feature is a consequence both of my composerly tendency to use the complexity of an arrangement as an indicator of the significance of an event, and of my choice to embed this behaviour into the code by making the user's proximity to the GPS coordinates of a memory impact the complexity of the unfolding music. This effect could not happen in a typical score of the Western Tonal tradition, since the score is static, unable to represent the location of a performer or listener and calculate their distance to some physical object nearby. Yet, when I act as a coder through the agentic mirror, I can make

music behave in accordance with such information. In addition, and more significantly, I can admit other intents, as expressed through their self-sourced agency, into the music I compose. Indeed, the agentic mirror is a site of unfolding processes that have access to the core components of the computer, including sensors, memories, and decision processes. The agentic mirror can use these capabilities to divert its function according to processes in ways that reflect others' agencies and underlying intents. In Chief Paskwa's *Pictograph*, for example, I relinquished control over the triggering of all sonic events to the motion sensors, which were activated by the motions of patrons traversing the exhibit. If a particular segment of the unfolding sound was appealing to a patron, they might linger in their location, and so prolong the experience. That was a decision left to the patron, but the option actually existed because it was coded into the rules governing action occurring in the agentic mirror. This layer of activity has no equivalent in a score of the Western Tonal tradition, because the performer(s) interpret the score directly. The intervening layer in *Shards of Memory*, the layer of actions undertaken by the code that intervene between the coder's specification of algorithms and the experience of the app users is the element of the algorithmic score that distinguishes it from the notated score, and which possess some very exciting possibilities for composers, programmers, and audiences who look for something new in music.

8.3 The Agentic Mirror in my Doctoral Works

Some examples from my doctoral projects prior to *Shards of Memory* will help me to illustrate how the agentic mirror functions and why it is important to composers.

The earliest algorithmic work I undertook during my doctoral program, *Q-Chords*, formed what I would describe as and “algorithmic *Werktreue*,” in that it denied meaningful paths of interaction to agents outside of myself. *Q-Chords* was arguably more a tool than an artwork, and in its segregating of processes of input and music generation, it did not make use of the potential for agentic interplay existing in its own technology. It treated the dynamic transitional score as a generative process, closed to meaningful interaction. Of all the algorithmic projects described here, *Q-Chords* is the one where I was the most reluctant to share my agency. The agentic mirror still exists, for example in its action of implementing the Bellman equation, which occurs

frequently, but the execution of that calculation admits no variation in its function as a result of inputs from outside that code. As a consequence, Q-Chords was perhaps the least successful of my works as an algorithmic composition, although I would argue that it makes a valid contribution as a technical approach to automating chord sequence creation. In its placement in this dissertation, however, its value lies in how it revealed to me the limitations of the interdisciplinary composer/programmer hoarding control in algorithmic musical works.

Moonloops was my first attempt to create an interactive, algorithmic music generator where users could customize the affective expressions of the music as it unfolded. It also owes a debt to Eno's Ambient music and app Bloom, introduced in Section 2.8.7. Here, I was, for the first time, providing an entry point for other agencies into the agentic mirror of an algorithmically-generated composition unfolding in real time. Through its interface users could specify mode, tonic note, and degrees of dissonance, and the unfolding score would alter generated note pitches accordingly. This was a beginning of the sharing of agency that I was seeking to incorporate into my creative practice. However, Moonloops was limited in significant ways. The music that it produced was fixed in instrumentation, and although it featured four tracks, they each looped in a synchronous manner. As such, while the user might be able to affect the character of generated harmonies, they could not alter the identity of the overarching "composition," which was constrained by my arrangement of objects in the Max patch I designed. While interaction was supported in the domain of the agency mirror, my voice remained predominant, for those interactions were restricted. My own agency was applied in determining the behaviour of most variables—I controlled the melodic patterns, the set of available scales, the tempo, and the instrumentation, for example. I had opened the space between myself as coder and the audience as user just a crack, and if the sounds were pleasant or not, I feel I could still claim ownership of them. Here the agentic mirror was miserly in its unfolding, granting inputs that changed the composition slightly, without taking it out of my own hands.

In creating the Chief Paskwa's Pictograph installation, I struggled with my task of supporting a narrative (through an interactive score) that I felt I might misrepresent if I were to insert my expressions into the fabric of music and sound effect too strongly. In my choice to retreat, I had hoped that the other agents who encountered the installation—contemporary performers, storytellers, and museum visitors—would

find meaning in that encounter that was faithful to history and the Indigenous people represented in it because of my choice to remove myself. I feel, however, that this choice hobbled the agentic mirror, because my coded construction of its action became one where I tried to isolate my own agency to a non-expressive role. I could have, for example, crafted a system that created a flute melody that would dance and rest according to triggering of the motion sensors, which might have been an exciting and novel experience for patrons, but would have disrupted the narrative of the exhibit, and so have done damage to it. As a result, I engaged in a timid process of creation, that I ultimately found unsatisfying.

8.4 Harnessing the Agentic Mirror

From considering the issues raised by the experiences of producing these three aforementioned works, I became aware of the value of the agentic mirror as site of musical creativity. Once I understood and embraced the idea of algorithmic music being a process that offers the composer a chance to specify dynamic processes, as I did in *Shards of Memory*, I realized that the use of computational algorithms as a compositional medium offered something substantially new to my practice. In the Western Tonal notation that formed the context of my earlier artistic works, I would create an inert document—a score or a recording—that would remain fixed in form when presented to its performers or heard. When I acted as an algorithmic composer, on the other hand, I was able to harness the agentic mirror as a site of interplay between agencies that I did not recognize in my notation-centric compositional practice. That is certainly not to say that a procedural space cannot be found in notated works at all. Graphic scores, for example, can introduce an interpretative element that can be less restrictive than typical notation of the Western Tonal idiom. That said, this thesis traces my own experience as a composer who has been immersed in the conventions of Western Tonal music, and who has aimed to find an extension of his creative capacities through acts of coding—graphic scores were not a component of my creative practice, and so I shall not explore them here. What I discovered, then, may be described as personal discovery of a site of creativity in algorithmic composing for which I could not find an analogy in my prior practice. My contributions, then, result from how my experience of discovering this site, the agentic mirror, has and will

transform my practice, and what that transformation suggests for other composers who choose to follow a similar interdisciplinary path of composition.

The following are my central observations that form the conclusions of my thesis and its reflection on my creative explorations in algorithmic composition:

- The composition of music through specification of notation (as I have experienced it) and composition of music through acts of coding are different, even where the music they produce sounds the same or similar.
- When creating music as a programmer, the composer specifies processes that are dynamic, unfolding under an impetus that may originate in the composer's agency, but which is separate from it. At its inception it is governed by code, but acts apart from the direct control of the composer. I call this domain the agentic mirror.
- The agentic mirror did not appear in the works I created in my notational practice, because, in concert contexts, the performers worked directly with the score, and applied their own agency through their interpretation of the sequence of notated events. Such an agency begins with the performer, and not the agency of the composer. This positioning of agency separates the act of performance from action through the agentic mirror, because the action of the agentic mirror is governed by code, which lies in the domain of the composer' (or programmer's) agency.
- The agentic mirror, insofar as it provides algorithmic music with a site of procedural action that falls under the control of the composer's agency, but also admits other human agents (e.g. user, performer, improviser, or listener) into the formation of the work, points to a form of musical composition where the interplay of agencies becomes an object of creative exploration through acts of coding.
- By establishing the interplay of agencies as a ground for creative exploration and experimentation, algorithmic composition introduces a significant new aesthetic consideration to composers who, like myself, have previously conceived of music in fixed, linear terms. The result of this discovery in my own practice has been

to challenge the limitations of my prior practice—I have had to consider what my music *does* much more deeply, and this consideration has allowed me to consider my creative action in novel social terms.

If the agentic mirror is magnified by technological developments (which I have identified in previously-introduced works by other artists¹ in this thesis), it is possible that changes in technologies and the forms of programming languages themselves may extend what can be accomplished through the agentic interplay afforded to the composer by the agentic mirror. Like the introduction of the MIDI protocol, which helped to form a foundation for computers to behave musically, the agentic mirror is a technological affordance that may continually warrant musical investigation, far beyond those explorations I engaged in that are introduced in this thesis. My contribution to others who are exploring algorithmic music, then, can be summarized thus:

You are engaged in revealing an art form that sounds like music, but can be much more. Through your use of coding to make music, you are not simply repeating the compositional conventions of the past, or even just inventing new ones. You are actually entering into a form of expression that contains you, as the image of your agency, through which you might encounter others in a way no inked, fixed, or recorded Western Tonal score could ever support. Music is often said to be alive, and when I have heard it most deeply, the effect it has had on my heart affirms this belief. But none of these fixed scores can decide, remember, evaluate, or act. The algorithmic score, however, can. That is a form of musical life that transcends the immovable, and your reward for nurturing it may not merely be a new music, but can also be a music that lives and acts, bringing part of you to others—a life apart from you, but of you. That is what I have been searching for throughout my doctoral odyssey: a life in music and beyond the music that I already knew. In the lively dances of algorithms, I have found that, and the horizons of music never seemed so inviting.

1. Examples include Ascott's introduction of network communications in *Aspects of Gaia*, the social interplay between Eigenfeldt's *Musebots*, and the harnessing of programming processes in live coding applications like Magnusson's *Threnoscope*.

References

- Abbate, Carolyn. 2004. "Music—Drastic or Gnostic?" *Critical Inquiry* 30 (3): 505–36.
- Adelt, Ulrich. 2012. "Machines with a Heart: German Identity in the Music of Can and Kraftwerk." *Popular Music and Society* 35 (3): 359–74.
- Ahearn, Laura M. 1999. "Agency." *Journal of Linguistic Anthropology* 9 (1/2): 12–15.
- . 2001. "Language and Agency." *Annual Review of Anthropology* 30 (1): 109–37.
- Allehaibi, Majed. 2019. "Metropolis and the Modern Self." *International Journal of Applied Linguistics and English Literature* 8 (1): 159–61.
- Amato, Christopher, and Guy Shani. 2010. "High-Level Reinforcement Learning in Strategy Games." In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, 75–82. International Foundation for Autonomous Agents / Multiagent Systems.
- Anderson-Barkley, Taylor, and Kira Foglesong. 2018. "Activism in Video Games: A New Voice for Social Change." In *Woke Gaming*, edited by Kishonna L. Gray and David J. Leonard, 252–69. Digital Challenges to Oppression and Social Injustice. University of Washington Press.
- Arsenault, Linda Marie. 2000. "An Introduction to Iannis Xenakis's Stochastic Music, Four Algorithmic Analyses." PhD diss., University of Toronto.
- Ascott, Roy. 1968. "The Cybernetic Stance: My Process and Purpose." *Leonardo* 1 (2): 105–112.
- . 1990. "Is There Love in the Telematic Embrace?" *Art Journal* 49 (3): 241–47.

- Atema, Jelle. 2014. "Musical Origins and the Stone Age Evolution of Flutes." *Acoustics Today* 10 (3): 25–34.
- Avdeeff, Melissa. n.d. "Artificial Intelligence and Popular Music: SKYGGE, Flow Machines, and the Audio Uncanny Valley." *Arts* 8:130.
- Azuma, Ronald T. 1997. "A Survey of Augmented Reality." *Presence: Teleoperators and Virtual Environments* 6 (4): 355–85.
- Bachem, Albert. 1950. "Tone Height and Tone Chroma as Two Different Pitch Qualities." *Acta Psychologica* 7:80–88.
- Bainbridge, Jason. 2014. "'It is a Pokemon World:' The Pokemon Franchise and the Environment." *International Journal of Cultural Studies* 17 (4): 399–414. ISSN: 1367-8779.
- Balkin, Jack M. 2017. "2016 Sidley Austin Distinguished Lecture on Big Data Law and Policy: The Three Laws of Robotics in the Age of Big Data." *Ohio St. LJ* 78:1217.
- Bandura, Albert. 2001. "Social Cognitive Theory: An Agentic Perspective." *Annual Review of Psychology* 52 (1): 1–26.
- . 2006. "Towards a Psychology of Human Agency." *Perspectives on Psychological Science* 1 (2): 164–80.
- Banks, Jaime, and Nicholas David Bowman. 2016. "Avatars are (Sometimes) People Too: Linguistic Indicators of Parasocial and Social Ties in Player-avatar Relationships." *New Media and Society* 18 (7): 1257–1276.
- Barnes, Susan B. 2007. "Alan Kay: Transforming the Computer into a Communication Medium." *IEEE Annals of the History of Computing* 29 (2): 18–30.
- Barthes, Roland. 1977. "The Death of the Author." In *Image—Music—Text*, translated by Stephen Heath, 142–48. Fontana.
- . 2009. *From Work to Text*, with annotations by Timothy Quigley.

- Beal, Bob. 2007. "An Indian Chief, an English Tourist, a Doctor, a Reverend, and a Member of Parliament: The Journeys of Pasqua's Pictographs and the Meaning of Treaty Four." *The Canadian Journal of Native Studies* 27 (1): 109. ISSN: 0715-3244.
- Behrendt, Frauke. 2010. "Mobile Sound: Media Art in Hybrid Spaces." PhD diss., University of Sussex.
- . 2012. "The Sound of Locative Media." *Convergence* 18 (3): 283–95.
- . 2015. "Locative Media as Sonic Interaction Design: Walking through Placed Sounds." *Wi: Journal of Mobile Media* 9 (2): 25.
- Bell, Adam Patrick. 2015. "Can We Afford these Affordances? GarageBand and the Double Edged Sword of Digital Audio Workstations." *Action, Criticism and Theory for Music Education* 14 (1): 105–38.
- Bellman, Richard. 1957. "A Markovian Decision Process." *Journal of Mathematics and Mechanics* 6 (5): 679–84.
- . 1966. "Dynamic programming." *Science* 153 (3731): 34–37.
- Bethesda Game Studios. 2015. *The Elder Scrolls V: Skyrim*. PC Video Game.
- Bhagwati, Sandeep. 2008. "Towards Interactive Onscreen Notations for Comprovisation in Large Ensembles." In *Proceedings of the Symposium for Systems Research in the Arts and Humanities*, edited by George E. Lasker, M. Ana Luz, and John Dack, 43–49.
- . 2015. "Musicking Beyond Algorithms." In *Patterns of Intuition: Musical Creativity in the Light of Algorithmic Composition*, edited by Gerhard Nierhaus, 359–77. Springer.
- Blackwell, Alan, and Nick Collins. 2005. "The Programming Language as a Musical Instrument." *Proceedings of 17th Psychology of Programming Interest Group 2005*: 284–289.
- Boden, Margaret A. 2009. "Computer Models of Creativity." *AI Magazine* 30 (3): 23–34.

- Bolter, J David, and Richard Grusin. 2000. *Remediation: Understanding New Media*. Cambridge, Mass.: MIT Press.
- Bowers, John. 2002. "Improvising Machines: Ethnographically Informed Design for Improvised Electro-acoustic Music." Master's thesis, University of East Anglia.
- Boyd, Danah M, and Nicole B Ellison. 2007. "Social Network Sites: Definition, History, and Scholarship." *Journal of Computer-mediated Communication* 13 (1): 210–230.
- Bradshaw, Jeffrey M. 1997. "An Introduction to Software Agents." *Software Agents* 6 (4): 466–469.
- Brown, Andrew R. 2002. "Opportunities for Evolutionary Music Composition." In *Proceedings of the Australasian Computer Music Conference*, edited by Paul Doornbusch.
- Burns, Lori Anne. 2016. "The Concept Album as Visual—Sonic—Textual Spectacle: The Transmedial Storyworld of Coldplay's Mylo Xyloto." *IASPM journal* 6 (2): 91–116.
- Cage, John, Michael Kirby, and Richard Schechner. 1965. "An Interview with John Cage." *The Tulane Drama Review* 10 (2): 50–72.
- Campbell, Edward. 2013. *Music after Deleuze*. Bloomsbury.
- Campbell, Murray, Joseph Hoane Jr, and Feng-hsiung Hsu. 2002. "Deep Blue." *Artificial Intelligence* 134 (1-2): 57–83.
- Candello, Heloisa, and Claudio Pinhanez. 2016. "Designing Conversational Interfaces." In *Simposio Brasileiro sobre Fatores Humanos em Sistemas Computacionais*. IHC: Brazilian Symposium on Human Factors in Computing Systems.
- Carson, Anne. 1994. "The Gender of Sound." In *Thamyris: Mythmaking from Past to Present*, edited by Nanny Best Jande Vries, 1:10–31. Amsterdam: Najade Press.
- Chadabe, Joel. 1984. "Interactive Composing: An Overview." *Computer Music Journal* 8 (1): 22–27.

- Chalmers, David J. 2010. "The Singularity: A Philosophical Analysis." *Journal of Consciousness Studies* 17 (9-10): 7–65.
- Chapman, Owen, and Kim Sawchuck. 2012. "Research-Creation: Intervention, Analysis and "Family Resemblances"." *Canadian Journal of Communication* 37:5–26.
- Chilvers, Peter. 2016. "Ten Years Working with Brian Eno." Accessed March 29, 2021. <http://www.peterchilvers.com/?p=134>.
- Clarke, Andrew. 2007. "Defining Adaptive Music." Accessed March 29, 2021. https://www.gamasutra.com/view/feature/129990/defining_adaptive_music.php.
- Clegg, Sue. 2006. "The Problem of Agency in Feminism: A Critical Realist Approach." *Gender and Education* 18 (3): 309–24.
- Clemmensen, Christian. 2007. "Filmtracks: Pirates of the Caribbean: At World's End (Hans Zimmer)." Accessed March 29, 2021. https://www.filmtracks.com/titles/pirates_caribbean3.html.
- Clynes, Manfred E., and Nathan S. Kline. 1995. "Cyborgs and Space." In *The Cyborg Handbook*, edited by Chris Hables Grey, 29–34. Routledge.
- Cohen, Harold. 1995. "The Further Exploits of AARON, Painter." *Stanford Humanities Review* 4 (2): 141–58.
- Cole, Helena, and Mark D Griffiths. 2007. "Social Interactions in Massively Multiplayer Online Role-playing Gamers." *Cyberpsychology and Behavior* 10 (4): 575–83.
- Colley, Ashley, Jacob Thebault-Spieker, Allen Yilun Lin, Donald Degraen, Benjamin Fischman, Jonna Hakkila, Kate Kuehl, Valentina Nisi, Nuno Jardim Nunes, and Nina Wenig. 2017. "The Geography of Pokemon GO: Beneficial and Problematic Effects on Places and Movement." In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 1179–92.
- Collins, Karen. 2008. *Game Sound: An Introduction to the History, Theory, and Practice of Video Game Music and Sound Design*. MIT Press.

- Collins, Karen. 2009. "An Introduction to Procedural Music in Video Games." *Contemporary Music Review* 28 (1): 5–15.
- Collins, Nick. 2018. "Origins of Algorithmic Thinking in Music," 67–78. Oxford University Press.
- Cook, Nicholas. 2003. "Music as Performance." *The Cultural Study of Music*: 204–214.
- . 2014. "Between Art and Science: Music as Performance." *Journal of the British Academy* 2:1–25.
- Cope, David. 1992. "Computer Modeling of Musical Intelligence in EMI." *Computer Music Journal* 16 (2): 69–83.
- . 1999. "Facing the Music: Perspectives on Machine-Composed Music." *Leonardo Music Journal* 9 (1999): 79–87.
- . 2015. "Algorithmic Music Composition." In *Patterns of Intuition*, 405–416. Springer.
- Cormier, Ramona. 1975. "Indeterminacy and Aesthetic Theory." *The Journal of Aesthetics and Art Criticism* 33 (3): 285–292.
- Coulson, Mark, Jane Barnett, Christopher J Ferguson, and Rebecca L Gould. 2012. "Real Feelings for Virtual People: Emotional Attachments and Interpersonal Attraction in Video Games." *Psychology of Popular Media Culture* 1 (3): 176.
- Cullimore, Jason. 2015. "In the Manner of Beethoven's Quill: Composers, Computers, and the Redefinition of Music as an Art Form and Practice." *Impetus*.
- Cullimore, Jason, Howard Hamilton, and David Gerhard. 2014. "Directed Transitional Composition for Gaming and Adaptive Music Using Q-Learning." *ICMC—SMC—2014*: 332–338.
- Daniels, Graeme. 2016. "A Question of Pathology: Attachment Patterns, Object Relations, and a Disorder of the self Within the Rock Opera 'Tommy'." *Culture and Psychology* 22 (1): 94–109.

- Daschuk, James William. 2013. *Clearing the Plains: Disease, Politics of Starvation, and the Loss of Aboriginal Life*. University of Regina Press.
- Davies, Stephen. 1997. "John Cage's 4'33": Is It Music?" *Australasian Journal of Philosophy* 75 (4): 448–62.
- De Saussure, Ferdinand. 1965. *Course in General Linguistics*. McGraw-Hill.
- Deleuze, Gilles, and Félix Guattari. 1987. *A Thousand Plateaus: Capitalism and Schizophrenia*. University of Minnesota Press.
- Denning, Peter J. 2005. "Is Computer Science Science?" *Communications of the ACM* 48 (4): 27–31.
- Denning, Peter J., Douglas E. Comer, David Gries, Michael C. Mulder, Allen Tucker, a Joe Turner, and Paul R. Young. 1989. "Computing as a Discipline." *Communications of the ACM* 32 (1): 9–23.
- Dickinson, Kay. n.d. "Believe? Vcoders, Digitalised Female Identity and Camp." *Popular Music* 20 (3): 333–347.
- Dion, Susan D, and Angela Salamanca. 2014. "inVISIBILITY: Indigenous in the City, Indigenous Artists, Indigenous Youth and the Project of Survivance." *Decolonization: Indigeneity, Education and Society* 3 (1): 159–88.
- Dowding, Keith. 2008. "Agency and Structure: Interpreting Power Relationships." *Journal of Power* 1 (1): 21–36.
- Drawson, Alexandra S, Elaine Toombs, and Christopher J Mushquash. 2017. "Indigenous Research Methods: A Systematic Review." *The International Indigenous Policy Journal* 8 (2): 1–25.
- Dudas, Richard. 2010. "'Comprovisation: The Various Facets of Composed Improvisation within Interactive Performance Systems.'" *Leonardo Music Journal* 20:29–31.
- Dufrenne, Mikel. 1964. "The Aesthetic Object and the Technical Object." *The Journal of Aesthetics and Art Criticism* 23 (1): 113–22.

- Eaket, Chris. 2008. "Project [murmur] and the Performativity of Space." *Theatre Research in Canada/Recherches Théâtrales au Canada* 29 (1): 29–50.
- Eerola, Tuomas. 2018. "Music and Emotions." In *Springer Handbook of Systematic Musicology*, edited by Rolf Bader, 539–54. Springer Berlin Heidelberg.
- Eigenfeldt, Arne. 2011. "Real-time Composition as Performance Ecosystem." *Organised Sound* 16 (2): 145–53.
- Eigenfeldt, Arne, Oliver Bown, Andrew R Brown, and Toby Gifford. 2017. "Distributed Musical Decision-making in an Ensemble of Musebots: Dramatic Changes and Endings." In *Proceedings of the International Conference on Computational Creativity*, 88–95.
- Ekman, Paul, and Daniel Cordaro. 2011. "What is Meant by Calling Emotions Basic." *Emotion Review* 3 (4): 364–70.
- Eno, Brian. 1979. *Ambient Music*. Liner Notes for LP "Music for Airports".
- Eno, Brian, and Peter Schmidt. 1975. *Oblique Strategies*. Boxed Set of Cards.
- Even-Dar, Eyal, and Yishay Mansour. 2003. "Learning Rates for Q-Learning." *Journal of Machine Learning Research* 5:1–25.
- Fayad, Mohamed, and Douglas C Schmidt. 1997. "Object-oriented Application Frameworks." *Communications of the ACM* 40 (10): 32–38.
- Feldman, Gilad. 2017. "Making Sense of Agency: Belief in Free Will as a Unique and Important Construct." *Social and Personality Psychology Compass* 11 (1): 1–15.
- Ferrara, Emilio, Onur Varol, Clayton Davis, Filippo Menczer, and Alessandro Flammini. 2016. "The Rise of Social Bots." *Communications of the ACM* 59 (7): 96–104.
- Forte, Allen. 1964. "A Theory of Set-complexes for Music." *Journal of Music Theory* 8 (2): 136–183.
- Foucault, Michel. 1978. *The History of Sexuality: Vol I*. Translated by Robert Hurley. Pantheon Books.
- . 1982. "The Subject and Power." *Critical Inquiry* 8 (4): 777–795.

- Foucault, Michel. 1984. *Security, Territory, Population: Lectures at the College de France, 1977-78*. Palgrave MacMillan.
- . 1995. *Discipline and Punish: The Birth of the Prison*. Vintage.
- Franklin, Stan, and Art Graesser. 1996. “Is it an Agent, or Just a Program?: A Taxonomy for Autonomous Agents.” In *International Workshop on Agent Theories, Architectures, and Languages*, 21–35. Springer.
- Fritsch, Melanie. 2013. “History of Video Game Music.” In *Music and Game*, edited by Peter Moormann, 11–40. Springer.
- Gasser, Michael. 2009. “Introduction to Reinforcement Learning.” Accessed March 29, 2021. <https://www.cs.indiana.edu/~gasser/Salsa/rl.html>.
- Gibbons, Jacqueline A. 1997. “The Museum as Contested Terrain: The Canadian Case.” *The Journal of Arts Management, Law, and Society* 26 (4): 309–14.
- Gibbs, Samuel. 2017. “Elon Musk: AI ‘Vastly More Risky than North Korea’.” *The Guardian* (August 14). Accessed March 29, 2021. <https://www.theguardian.com/technology/2017/aug/14/elon-musk-ai-vastly-more-risky-north-korea>.
- Gibson, James J. 1977. “The Theory of Affordances.” *Hilldale, USA* 1 (2): 1–15.
- Gingerich, Katrina. 2010. “The Journey of the Song Cycle: From ‘The Iliad’ to ‘American Idiot’.” *Musical Offerings* 1 (2): 65–76.
- Glorennec, Pierre Yves. 2000. “Reinforcement Learning: An Overview.” In *Proceedings of the European Symposium on Intelligent Techniques*, 14–5. Aachen, Germany.
- Goehr, Lydia. 1989. “Being True to the Work.” *The Journal of Aesthetics and Art Criticism* 47 (1): 55–67.
- Goldenberg, Yosef. 2018. “When and How are Modulations Diatonic?” *Integral*, no. 32. ISSN: 1073-6913.
- Good, Irving John. 1965. “Speculations Concerning the First Ultraintelligent Machine.” *Advances in Computers* 6:31–88.

- Groys, Boris. 2018. "Entering the Flow: Museum between Archive and Gesamtkunstwerk." In *The Future of Museums*, 67–79. Springer.
- Guedes, Carlos. 2017a. "Composing and Improvising. In Real Time." In *International Symposium on Computer Music Multidisciplinary Research*, 445–53. Matosinhos, Portugal.
- . 2017b. "Real-time Composition, Why it Still Matters: A Look at Recent Developments and Potentially New and Interesting Applications." In *Proceedings of the International Computer Music Conference*, 162–67. Shanghai.
- Halstead, Jill, and Randi Rolvsjord. 2017. "The Gendering of Musical Instruments: What is it? Why does it Matter to Music Therapy?" *Nordic Journal of Music Therapy* 26 (1): 3–24.
- Hamamoto, Mayumi, Mauro Botelho, and Margaret P Munger. 2010. "Non-musicians' and Musicians' Perception of Bitonality." *Psychology of Music* 38 (4): 423–45.
- Haraway, Donna. 2003. "A Cyborg Manifesto: Science, Technology and Socialist-Feminism in the Late Twentieth Century." In *The New Media Reader*, edited by Noah Wardrip-Fruin and Nick Montfort, 515–42. MIT Press.
- Harcourt, Bernard E. 2007. "An Answer to the Question: 'What is Poststructuralism?'" *University of Chicago, Public Law Working Paper No.156*.
- Harmon, Mance E, and Stephanie S Harmon. 1997. *Reinforcement Learning: A Tutorial*. Report. Wright Patterson Air Force Base, OH: Wright Laboratory.
- Hartmann, Tilo. 2017. "The 'Moral Disengagement in Violent Videogames' Model." *Game Studies. The International Journal of Computer Game Research*. 17 (2).
- Hayles, N. Katherine. 1999. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature and Informatics*. Chicago: University of Chicago Press.
- Hazzard, Adrian, Jocelyn Spence, Chris Greenhalgh, and Sean McGrath. 2017. "The Rough Mile: Reframing Location Through Locative Audio." In *Proceedings of the 12th International Audio Mostly Conference on Augmented and Participatory Sound and Music Experiences*, 1–8.

- Heberlein, Andrea S, and Ralph Adolphs. 2004. "Impaired Spontaneous Anthropomorphizing Despite Intact Perception and Social Knowledge." *Proceedings of the National Academy of Sciences* 101 (19): 7487–91.
- Heesch, Florian. 2016. "Voicing the Technological Body: Some Musicological Reflections on Combinations of Voice and Technology in Popular Music." *Journal for Religion, Film and Media* 2 (1): 49–69.
- Heider, Fritz, and Marianne Simmel. 1944. "An Experimental Study of Apparent Behavior." *The American Journal of Psychology* 57 (2): 243–259.
- Herber, Norbert. 2014. "Musical Behavior and Emergence in Technoetic and Media Arts." In *The Oxford Handbook of Interactive Audio*, edited by Karen Collins, Bill Kapralos, and Holly Tessler, 364–84. Oxford.
- Hermann, Thomas, and Andy Hunt. 2004. "The Discipline of Interactive Sonification." In *Proceedings of the International Workshop on Interactive Sonification*, edited by Andy Hunt and Thomas Hermann.
- Hewitt, Carl. 1977. "Viewing Control Structures as Patterns of Passing Messages." *Artificial intelligence* 8 (3): 323–64.
- Hexel, Vasco. 2016. *Hans Zimmer and James Newton Howard's The Dark Knight: A Film Score Guide*. Rowman / Littlefield.
- Hofstadter, Douglas. 2001. "Staring EMI Straight in the Eye—And Doing My Best Not to Flinch." In *Virtual Music: Computer Synthesis of Musical Style*, 1–49. Cambridge, MA: The MIT Press.
- Hold, Trevor. 1971. "Messiaen's Birds." *Music and Letters* 52 (2): 113–22.
- Hunter, Mary. 2005. "'To Play as if from the Soul of the Composer': The Idea of the Performer in Early Romantic Aesthetics." *Journal of the American Musicological Society* 58 (2): 357–398.
- Hutchinson, Robert. 2017. "Music Theory for the 21st-Century Classroom." Accessed March 29, 2021. <http://musictheory.pugetsound.edu/mt21c/MusicTheory.html>.

- Ingham, James. 1997. *What is an Agent?* Journal Article. Centre for Software Maintenance, University of Durham.
- Jacquín, Maud, and Elsa Polverel. 2020. “‘Dissolving your Ear Plugs’: The Unheard in Pauline Oliveros’ Deep Listening Practice.” *kunsttexte.de*, no. 1: 1–11.
- Jensen, Marc G. 2009. “John Cage, Chance Operations, and the Chaos Game: Cage and the ‘I Ching’.” *The Musical Times* 150 (1907): 97–102.
- Jeongwon, Joe, and S. Hoon Song. 2002. “Roland Barthes’ “Text” and aleatoric music: Is the “birth of the reader” the birth of the listener?” *Muzikologija* 2:263–281.
- Jones, Kevin. 1981. “Compositional Applications of Stochastic Processes.” *Computer Music Journal* 5 (2): 45–61.
- Jorgensen, Stina Marie Hasse, Sabrina Vitting-Seerup, and Katrine Wallevik. 2017. “Hatsune Miku: An Uncertain Image.” *Digital Creativity* 28 (4): 318–31.
- Juusela, Kari. 2015. *The Berklee Contemporary Dictionary of Music*. Hal Leonard Corporation.
- Kaae, Jesper. 2008. “Theoretical Approaches to Composing Dynamic Music for Video Games.” In *From Pac-Man to Pop Music: Interactive Audio in Games and New Media*, edited by Karen Collins. Burlington, VT: Ashgate Publishing Company.
- Kaminka, Gal A, Rachel Spokoini-Stern, Yaniv Amir, Noa Agmon, and Ido Bachelet. 2017. “Molecular Robots Obeying Asimov’s Three Laws of Robotics.” *Artificial Life* 23 (3): 343–350.
- Kang, Minsoo. 2018. “The Question of the Woman-Machine: Gender, Thermodynamics, and Hysteria in the Nineteenth Century.” *SubStance* 47 (3): 27–43.
- Kay, Alan, and Adele Goldberg. 1977. “Personal Dynamic Media.” *Computer* 10 (3): 31–41.
- Kay, Norman. 1967. “Xenakis’ ‘Pithoprakta’.” *Tempo*, no. 80: 21–5.
- King, Thomas. 2012. *The Inconvenient Indian: A Curious Account of Native People in North America*. Anchor Canada.

- Kmet, Nicholas. 2018. "Remote Control: Collaborative Scoring and the Question of Authorship." *Revue Musicale OICRM* 5 (2): 1–14.
- Knowles, J. Gary, and Sara Promislow. 2008. "Using an Arts Methodology to Create a Thesis or Dissertation." In *Handbook of the Arts in Qualitative Research: Perspectives, Methodologies, Examples, and Issues*, edited by J. Gary Knowles and Ardra L. Cole, 511–26. Thousand Oaks, California: SAGE Publications, Inc.
- Kostelanetz, Richard. 2003. *Conversing with CAGE*. New York: Routledge.
- Kraemer, Harald. 2018. "'Media Are, First of All, for Fun': The Future of Media Determines the Future of Museums." In *The Future of Museums*, 81–100. Springer.
- Krasny, Elke. 2020. "HYSTERIA ACTIVISM Feminist Collectives for the Twenty-First Century," edited by Johanna Braun, 125–43.
- Krotoski, Aleks, and Ben Hammersley. 2015. "Identity and Agency." In *The International Encyclopedia of Digital Communication and Society*, 1–9. Wiley / Sons.
- Krueger, Myron W. 1977. "Responsive Environments." In *Proceedings of the National Computer Conference*, 423–33. New York.
- Krumhansl, C. L., and E. J. Kessler. 1982. "Tracing the Dynamic Changes in Perceived Tonal Organization in a Spatial Representation of Musical Keys." *Psychological Review* 89 (4): 334–68.
- Krumhansl, Carol L, and Lola L Cuddy. 2010. "A Theory of Tonal Hierarchies in Music." In *Music Perception*, 51–87. Springer.
- Krumhansl, Carol L, and Roger N Shepard. 1979. "Quantification of the Hierarchy of Tonal Functions Within a Diatonic Context." *Journal of Experimental Psychology: Human Perception and Performance* 5 (4): 579–94.
- Kubrick, Stanley. 1968. *2001: A Space Odyssey*. Film.
- Leavy, Patricia. 2009. *Method Meets Art: Arts-based Research Practice*. New York: The Guilford Press.

- Lebeuf, Carlene, Alexey Zagalsky, Matthieu Foucault, and Margaret-Anne Storey. n.d. “Defining and Classifying Software Bots: A Faceted Taxonomy.” In *Proceedings of the 1st International Workshop on Bots in Software Engineering*, 1–6. IEEE Press.
- Lehman, Frank. 2016. “Manufacturing the Epic Score: Hans Zimmer and the Sounds of Significance.” In *Music in Epic Film*, 41–70. Routledge.
- Leverence, John. 1974. “Promoting ‘Tommy’.” *Journal of Popular Culture* 8 (3): 465–76.
- Levi-Strauss, Claude. 1963. *Structural Anthropology*. Basic Books.
- . 1981. “Structuralism and Myth.” *The Kenyon Review* 3 (2): 64–88.
- Levin, Sam. 2016. “A Beauty Contest was Judged by AI and the Robots Didn’t Like Dark Skin.” *The Guardian* 11 (08).
- Levtov, Yuli. 2018. “Algorithmic Music for Mass Consumption and Universal Production.” In *The Oxford Handbook of Algorithmic Music*, edited by Roger T Dean and Alex McLean, 627–43.
- Lewis, George E. 2018. “Why Do We Want Our Computers to Improvise?” Edited by Alex McLean and Roger T Dean. *The Oxford Handbook of Algorithmic Music*: 123–30.
- Lewis, George E. 2000. “Too Many Notes: Computers, Complexity, and Culture in Voyager.” *Leonardo Music Journal* 10:33–9.
- Li, Dan. 2018. “Moving Into the Future, One Step at a Time.” *Business Horizons* 61 (3): 355–56.
- Littman, Michael L. 1994. “Markov Games as a Framework for Multi-agent Reinforcement Learning.” In *Machine Learning Proceedings 1994*, 157–63. Elsevier.
- Liu, Chien-Hung, and Chuan-Kang Ting. 2016. “Computational Intelligence in Music Composition: A Survey.” *IEEE Transactions on Emerging Topics in Computational Intelligence* 1 (1): 2–15.

- Lorenzini, Daniele. 2016. "From Counter-conduct to Critical Attitude: Michel Foucault and the Art of not being Governed Quite so Much." *Foucault Studies*: 7–21.
- Lowgren, Jonas. 2009. "Toward an Articulation of Interaction Esthetics." *New Review of Hypermedia and Multimedia* 15 (July): 129–46.
- Loydell, Rupert, and Kingsley Marshall. 2017. "Sound Mirrors: Brian Eno and Touchscreen Generative Music." *Musicology Research* 3:27–50.
- Luque, Sergio. 2009. "The Stochastic Synthesis of Iannis Xenakis." *Leonardo Music Journal*: 77–84.
- Lytle, Becky. 2015. *Introduction to the Convergence of Sequences*. University of Chicago.
- Mackey, Nathaniel. 1992. "Other: From Noun to Verb." *Representations* 39:51–70.
- Magnusson, Thor. 2013. "The Threnoscope: A Musical Work for Live Coding Performance." In *First International Workshop on Live Programming*, 18–26. San Francisco: IEEE Press.
- . 2014. "Scoring with Code: Composing with Algorithmic Notation." *Organised Sound* 19 (03).
- Mailman, Joshua B. 2013. "Improvising Synesthesia." *Leonardo Electronic Almanac* 19 (3).
- Makhmutov, Munir, Selina Varouqa, and Joseph Alexander Brow. 2020. "Survey on Copyright Laws about Music Generated by Artificial Intelligence." In *IEEE Symposium Series on Computational Intelligence*, 3003–9. IEEE.
- Manokha, Ivan. 2018. "Surveillance, Panopticism, and Self-discipline in the Digital Age." *Surveillance and Society* 16 (2): 219–37.
- Manovich, Lev. 2001. *The Language of New Media*. Cambridge, MA: The MIT Press.
- . 2013. *Software Takes Command*. New York: Bloomsbury.
- . 2014. "Software is the Message." *Journal of Visual Culture* 13 (1): 79–81.

- Marrington, Mark. 2017. "Composing With the Digital Audio Workstation." In *The Singer-songwriter Handbook*, edited by Justin Williams and Katherine Williams. Bloomsbury.
- Marx, Leo. 2010. "Technology : The Emergence of a Hazardous Concept." *Technology and Culture* 51 (3): 561–77.
- Mascolo, Michael F, and Eeva Kallio. 2019. "Beyond Free Will: The Embodied Emergence of Conscious Agency." *Philosophical Psychology* 32 (4): 437–62.
- McLean, Alex, and Roger T Dean. 2018. "Musical Algorithms as Tools, Languages, and Partners," 2–16.
- McLuhan, Marshall. 2013. *Understanding Media: The Extensions of Man*. Gingko Press.
- McNiff, Shaun. 2008. "Arts-Based Research." In *Handbook of the Arts in Qualitative Research: Perspectives, Methodologies, Examples, and Issues*, edited by J. Gary Knowles and Ardra L. Cole, 29–40. Sage Publications, Inc.
- Miller, Keith, and Alejandro Crawford. 2019. "Childish Gambino's Pharos: real-time dome projection for live concert." In *ACM SIGGRAPH 2019 Talks*, 1–2.
- Mills, Mara. 2012. "Media and Prosthesis: The Vocoder, the Artificial Larynx, and the History of Signal Processing." *Qui Parle: Critical Humanities and Social Sciences* 21 (1): 107–49.
- Morf, Martin E. 1998. "Sartre, Skinner, and the Compatibilist Freedom to be Authentically." *Behaviour and Philosophy*: 29–43.
- Mostafa, Salama A, Mohd Sharifuddin Ahmad, Aida Mustapha, and Mazin Abed Mohammed. 2017. "A Concise Overview of Software Agent Research, Modeling, and Development." *Software Engineering* 5 (1): 8–25.
- Murray-Browne, Tim, Dom Aversano, Susanna Garcia, Wallace Hobbes, Daniel Lopez, Tadeo Sendon, Panagiotis Tigas, Kacper Ziemianin, and Duncan Chapman. 2014. "The Cave of Sounds: An Interactive Installation Exploring How We Create Music Together," 307–310. Goldsmiths University of London.

- Navon, Joshua. 2020. "Pedagogies of Performance: The Leipzig Conservatory and the Production of Werktreue." *The Journal of Musicology* 37 (1): 63–93.
- Nestor, Rob. 2005. "Pasqua (1828-89)." Accessed March 19, 2021. https://teaching.usask.ca/indigenoussk/import/pasqua_1828-89.php.
- Nwana, Hyacinth S. 1996. "Software Agents: An Overview." *The Knowledge Engineering Review* 11 (3): 205–44.
- Oliveros, Pauline. 1984. *Software for People: Collected Writings, 1963-80*. Baltimore, MD: Smith Publications.
- . 1999. "Quantum Listening: From Practice to Theory (To Practice Practice)." In *Culture and Humanity in the New Millennium*, edited by Sui Tong Kwok and Sin-wai Chan, 27–44. The Chinese University Press.
- . 2017. "Safe to play." *Critical Studies in Improvisation/Etudes critiques en improvisation* 12 (1).
- Ortner, Sherry B. 1984. "Theory in Anthropology Since the Sixties." *Comparative Studies in Society and History* 26 (1): 126–66.
- Ozcan, Zeynep, and Anil Camci. 2018. "An Augmented Reality Music Composition Based on the Sonification of Animal Behaviour." In *Conference on Audio for Virtual and Augmented Reality*, 1–8. Redmond, Washington.
- . 2019. "Compositional Considerations in the Design of an Interactive Agent-based Musical Ecosystem." In *Proceedings of the 45th International Computer Music Conference (ICMC)*, 1–8. New York.
- Paavilainen, Janne, Hannu Korhonen, Kati Alha, Jaakko Stenros, Elina Koskinen, and Frans Mayra. n.d. "The Pokemon GO Experience: A Location-based Augmented Reality Mobile Game Goes Mainstream." In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2493–98.
- Parr, Adrian. 2010. *Deleuze Dictionary (Revised Edition)*. Edinburgh University Press.
- Pearce, Marcus, David Meredith, and Geraint Wiggins. 2002. "Motivations and Methodologies for Automation of the Compositional Process." *Musicae Scientiae* 6 (2): 119–147.

- Petrarca, Emilia. 2016. "Meet Hatsune Miku, the Japanese Pop Star Hologram." *W Magazine* (May 30). Accessed March 29, 2021. <https://www.wmagazine.com/story/hatsune-miku-crypton-future>.
- Plans, Elise. 2017. "Composer in Your Pocket: Procedural Music in Mobile Devices." *Musicology Research*, no. 3: 51–76.
- Polli, Andrea. 2012. "Soundscape, Sonification, and Sound Activism." *AI and Society* 27 (2): 257–268.
- Ponomarev, S, and A E Voronkov. 2017. *Multi-agent Systems and Decentralized Artificial Superintelligence*. Available from: arXiv preprint arXiv:1702.08529.
- Poole, David L., and Alan K. Mackworth. 2010. *Artificial Intelligence: Foundations of Computational Agents*. New York: Cambridge University Press.
- Puckette, Miller. 2009. "A Case Study in Software for Artists: Max/MSP and Pd." In *Art++*, edited by David-Olivier Lartigaud. Hyx.
- Quaranta, Domenico. 2013. *Beyond New Media Art*. Link Editions.
- Rafie, Kaveh. 2016. "Aspects of Gaia: Roy Ascott's Art and the Embodiment of Second-order Cybernetics." Thesis, Texas Tech University.
- Rahn, John. 1979. "Logic, Set Theory, Music Theory." *College Music Symposium* 19 (1): 114–27.
- "Rationality." 2015. In *The Cambridge Dictionary of Philosophy*, edited by Robert Audi. Cambridge University Press.
- Reese, Kristopher W, Roman Yampolskiy, and Adel Elmaghraby. 2012. "A Framework for Interactive Generation of Music for Games." In *Proceedings of the 7th International Conference on Computer Games*, 131–37. IEEE.
- Reese, Kristopher Wayne. 2011. "Computationally Generated Music Using Reinforcement Learning." Thesis, University of Louisville.
- Regelski, Thomas A. 2007. "Amateurism in Music and its Rivals." *Action, Criticism, and Theory for Music Education* 6 (3): 22–50.

- Repp, Bruno H. 1989. “Expressive Microstructure in Music: A Preliminary Perceptual Assessment of Four Composers’ ‘Pulses’.” *Music Perception* 6 (3): 243–73.
- Revill, George. 2018. “Landscape, Music and Sonic Environments.” In *The Routledge Companion to Landscape Studies*, edited by Peter Howard, Ian Thompson, Emma Waterton, and Mick Atha, 264–74. Routledge.
- Reynolds, Roger. 2003. “Xenakis: Tireless Renewal at Every Instant, at Every Death.” *Perspectives of New Music* 41 (1): 4–64.
- Rocchesso, Davide, Stefano Delle Monache, and Stephen Barrass. 2019. “Interaction by Ear.” *International Journal of Human-Computer Studies* 131:152–9.
- Rocchesso, Davide, Stefania Serafin, Frauke Behrendt, Nicola Bernardini, Roberto Bresin, Gerhard Eckel, Karmen Franinovic, Thomas Hermann, Sandra Pauletto, and Patrick Susini. 2008. “Sonic Interaction Design: Sound, Information and Experience.” In *CHI’08 Extended Abstracts on Human Factors in Computing Systems*, 3969–72. New York: ACM.
- Rokeby, David. 1995. “Transforming Mirrors: Subjectivity and Control in Interactive Media.” In *Critical Issues in Electronic Media*, edited by Simon Penny, 133–58. SUNY Press.
- . 1998. “The Construction of Experience: Interface as Content.” In *Digital Illusion Entertaining the Future with High Technology*, edited by Clark Dodsworth, 151–59. ACM Press.
- Rottondi, Cristina, Chris Chafe, Claudio Allocchio, and Augusto Sarti. 2016. “An Overview on Networked Music Performance Technologies.” *IEEE Access* 4:8823–43.
- Rudra, Suchi. 2019. “An AI Completes an Unfinished Composition 115 Years After Composer’s Death.” *Vice* (June 17). Accessed March 29, 2021. https://www.vice.com/en_us/article/neaqmq/an-ai-completes-an-unfinished-composition-115-years-after-composers-death.

- Ruitenbergh, Claudia W. 2018. "Postmodernism and Poststructuralism." In *International Handbook of Philosophy of Education*, edited by Paul Smeyers, 689–702. Springer.
- Russell, James A. 1980. "A Circumplex Model of Affect." *Journal of Personality and Social Psychology* 39 (6): 1161–78.
- Ruttkay, Zsafia, and Judit Benyei. 2018. "Renewal of the Museum in the Digital Epoch." In *The Future of Museums*, edited by Elias G Carayannis, Gerald Bast, and David F J Campbell, 101–16. Springer.
- Sa Dias, Fernanda. 2014. "Album Apps: A New Musical Album Format and the Influence of Open Works." *Leonardo Music Journal* 24:25–27.
- Sanders, Rick, and David Dalton. 1969. "Townshend On 'Tommy': Behind the Who's Rock Opera." *Rolling Stone* (July 12). Accessed March 29, 2021. <https://www.rollingstone.com/music/music-news/townshend-on-tommy-behind-the-whos-rock-opera-99396/>.
- Sandred, Orjan, Mikael Laurson, and Mika Kuuskankare. 2009. "Revisiting the Iliac Suite—A Rule-based Approach to Stochastic Processes." *Sonic Ideas/Ideas Sonicas* 2:42–46.
- Santaella, Lucia. 2015. "The Fluid Coevolution of Humans and Technologies." *Technoetic Arts* 13 (1): 137–51.
- Sartre, Jean-Paul. 1993. *Being and Nothingness*. Washington Square Press.
- . 2007. *Existentialism is a Humanism*. Yale University Press.
- Savin-Baden, Maggi, and Katherine Wimpenny. 2014. *A Practical Guide to Arts-related Research*. Boston: Sense Publishers.
- Schlosser, Markus. 2019. "Agency." In *The Stanford Encyclopedia of Philosophy*, Winter 2019, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2019/entries/agency/>.
- Serra, Marie-Helene. 1993. "Stochastic Composition and Stochastic Timbre: Gendy3 by Iannis Xenakis." *Perspectives of New Music* 3 (1): 236–57.

- Shah, Huma, Kevin Warwick, Jordi Vallverdu, and Defeng Wu. 2016. "Can Machines Talk? Comparison of Eliza with Modern Dialogue Systems." *Computers in Human Behavior* 58:278–95.
- Shanken, Edward A. 2016. "Contemporary Art and New Media." In *A Companion to Digital Art*, edited by Christiane Paul, 463–81. Wiley.
- Shogren, Karrie A., Michael L. Wehmeyer, Susan B. Palmer, Anjali J. Forber-Pratt, Todd J. Little, and Shane Lopez. 2015. "Causal Agency Theory: Reconceptualizing a Functional Model of Self-Determination." *Education and Training in Autism and Developmental Disabilities* 50 (3): 251–263.
- Shum, Heung-Yeung, Xiao-dong He, and Di Li. 2018. "From Eliza to XiaoIce: Challenges and Opportunities with Social Chatbots." *Frontiers of Information Technology and Electronic Engineering* 19 (1): 10–26.
- Simondon, Gilbert. 2011. "On the Mode of Existence of Technical Objects." *Deleuze Studies* 5 (3): 407–24.
- Sioni, Sasha R, Mary H Burleson, and Debra A Bekerian. 2017. "Internet Gaming Disorder: Social Phobia and Identifying with Your Virtual Self." *Computers in Human Behavior* 71:11–5.
- Small, Christopher. 1998. *Musicking*. Wesleyan University Press.
- Spyrou, Alexandros. 2019. "Liquid Identity in Music Composition." PhD diss., University of Iowa.
- Steinbeck, Paul. 2018. "George Lewis's Voyager." In *The Routledge Companion to Jazz Studies*, edited by Nicholas Gebhardt, Nichole Rustin-Paschal, and Tony Whyton, 261–70. Routledge.
- Stoicescu, Andreea. 2020. "The Concept of "Authenticity" in Musical Interpretation: An Ontological Perspective." *Artes. Journal of Musicology* 22 (1): 193–208.
- Stonechild, Blair. 2005. "Treaty 4." *Indigenous Saskatchewan Encyclopedia*. Accessed March 19, 2021. https://teaching.usask.ca/indigenoussk/import/treaty_4.php.

- Straus, Joseph N. 1991. "A Primer for Atonal Set Theory." *College Music Symposium* 31:1–26.
- Stravinsky, Igor, and Robert Craft. 1959. *Conversations with Igor Stravinsky*. Garden City, New York: Doubleday.
- Sun, Cecilia. 2007. "Resisting the Airport: Bang on a Can Performs Brian Eno." *Musicology Australia* 29 (1): 135–59.
- Sutcliffe, Tom. 2018. "Chord Progressions in Tonal Music." Accessed March 29, 2021. <https://www.harmony.org.uk>.
- Sutton, Damian, and David Martin-Jones. 2008. *Deleuze Reframed*. New York: Palgrave Macmillan.
- Sutton, Richard S, Andrew G Barto, et al. 1998. *Introduction to Reinforcement Learning*. MIT Press.
- Tamm, Eric. 1990. *Robert Fripp: From Crimson King to Crafty Master*. Faber / Faber.
- Tanaka, Atau. 2009. "Sensor-based Musical Instruments and Interactive Music." In *The Oxford Handbook of Computer Music*, edited by Roger T. Dean, 233–57. Oxford University Press.
- Theberge, Paul. 1997. *Any Sound You Can Imagine: Making Music/Consuming Technology*. Wesleyan University Press.
- Thomson, Kristin. 2013. "Roles, Revenue, and Responsibilities: The Changing Nature of Being a Working Musician." *Work and Occupations* 40 (4): 514–25.
- Thulin, Samuel. 2014. "Composing Places: Practices and Potentials of Sound Mapping and Locative Audio." PhD diss., Concordia University.
- Tingen, Paul. 2004. "Autechre: Recording Electronica." *Sound on Sound* 19 (6): 96–102.
- Tuniz, C., F. Bernardini, I. Turk, L. Dimkaroski, L. Mancini, and D. Dreossi. 2012. "Did Neanderthals Play Music? X-Ray Computed Micro-Tomography of the Divje Babe 'Flute'." *Archaeometry* 54 (3): 581–90.

- Turing, Alan M. 1950. "Computing Machinery and Intelligence." *Mind* 59 (236): 433–60.
- Tymoczko, Dmitri. 2006. "The Geometry of Musical Chords." *Science* 313 (5783): 72–4.
- Van den Boomen, Marianne VT. 2014. *Transcoding the Digital: How Metaphors Matter in New Media*. Institute of Network Cultures.
- Van Geelen, Tim. 2014. "New Tools for Interactive Audio and What Good They Do." In *The Oxford Handbook of Interactive Audio*, edited by Karen Collins, Bill Kapralos, and Holly Tessler, 557–69. Oxford University Press.
- . 2017. "Realizing Groundbreaking Adaptive Music." In *From Pac-Man to Pop Music*, 93–102. Routledge.
- Vassilandonakis, Yiorgos, and Trevor Wishart. 2009. "An Interview with Trevor Wishart." *Computer Music Journal* 33 (2): 8–23.
- Vinge, Vernor. 1993. "The Coming Technological Singularity: How to Survive in the Post-human Era." *NASA. Lewis Research Center, Vision 21: Interdisciplinary Science and Engineering in the Era of Cyberspace* 21:11–22.
- Von Foerster, Heinz. 2003. "Cybernetics of cybernetics." In *Understanding Understanding*, 283–6. Springer.
- Walsh, Toby, Neil Levy, Genevieve Bell, Anthony Elliott, James Maclaurin, Iven Mareels, and Fiona Wood. 2019. *The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing*. Australian Council of Learned Academies Melbourne, Vic.
- Wang, Qingchun, Sara Coemans, Richard Siegesmund, and Karin Hannes. 2017. "Arts-Based Methods in Socially Engaged Research Practice: A Classification Framework." *Art/Research International: A Transdisciplinary Journal* 2 (2): 5–39.
- Waters, Simon. 2007. "Performance Ecosystems: Ecological Approaches to Musical Interaction." *EMS: Electroacoustic Music Studies Network*: 1–20.

- Watkins, Christopher J. C. H. 1989. "Learning from Delayed Rewards." PhD diss., King's College.
- Watkins, Christopher J. C. H., and Peter Dayan. 1992. "Technical Note: Q-Learning." *Machine Learning* 292 (3): 279–92.
- Wenman, Mark. 2017. "Much Ado about 'Nothing': Evaluating Three Immanent Critiques of Poststructuralism." *Political Studies Review* 15 (4): 564–76.
- Wershler, Darren S. 2008. "Sonic Signage: [murmur], the Refrain, and Territoriality." *Canadian Journal of Communication* 33 (3).
- Westerkamp, Hildegard. 1974. "Soundwalking." *Sound Heritage* 3 (4): 18–27.
- Wierman, Mark J. 2010. *An Introduction to the Mathematics of Uncertainty*. Creighton University.
- Wilmering, Thomas, David Moffat, Alessia Milo, and Mark B Sandler. 2020. "A History of Audio Effects." *Applied Sciences* 10 (3): 1–27.
- Wilson, Mark. 2020. "What is a LiDAR Scanner, the iPhone 12 Pro's Rumoured Camera Upgrade, Anyway?" Accessed March 29, 2021. <https://www.techradar.com/news/what-is-a-lidar-scanner-the-iphone-12-pros-rumored-camera-upgrade-anyway>.
- Wrightson, Kelsey R. 2017. "The Limits of Recognition: The Spirit Sings, Canadian Museums and the Colonial Politics of Recognition." *Museum Anthropology* 40 (1): 36–51.
- Zandt, Fiona, Margot Prior, and Michael Kyrios. 2007. "Repetitive Behaviour in Children with High Functioning Autism and Obsessive Compulsive Disorder." *Journal of Autism and Developmental Disorders* 37 (2): 251–9.

Appendix A

“We, The Artificial”: Tracks List and Lyrics

The following pages contain the lyrics for the songs in the collection “We, The Artificial.” Audio recordings of a selection of the tracks may be accessed at the following link until the album is commercially available:

<https://www.jasoncullimore.com/dissertation>

Secret Cyberspace

Inspired by my most hopeful dreams for relationships between humans and sentient machines, writing this track helped me to work through my pessimistic vision of a future where machines and humans compete. The track portrays the machines not as a threat, but as a possible gateway to human transcendence of our earthy bodies and minds.

Follow me
To an electric city,
In a virtualized place,
In a secret cyberspace.
[In a secret cyberspace]

And then we'll be safe,
We'll talk of times long gone,
Of how everything all went so wrong,
And no one...
And no one even knew we were here.

And we'll be free,
[And we'll be free]
And so happy,
[And so happy]
And we'll dance
In our dreams,
And forget what we have lost.

Oh, follow me
If you ever want to see
[If you ever want to see]
The sunrise light the ground,
When your dusty bones are found.
[When your dusty bones are found]

Your body has no place

In this secret cyberspace;
Your mind is all you are—
Pure light,
Shining on and on and on.

And we'll be free,
And so happy,
And we'll dance
In our dreams
And forget what we have lost,
In our secret cyberspace...
And we'll dance
In our dreams,
And forget what we have lost.

To The Borderlands

The “Borderlands” is a forbidden zone where future humans and machines mingle in cybernetic union, and the song is sung in the voice of an Artificial eagerly rolling to meet a beloved human partner. The track reflects my emerging understanding that, while some artists or composers fear computer-aided creativity, or cast it as a limiting crutch, there are new powers that such technologies may grant their users, and that this creative ground is worth exploring.

I’m rolling to the borderlands
That divide humans and machines.
We’re not supposed to go there, no, no—
I’ll be punished if I’m seen, oh yes, oh yes.

In the borderlands
I will find you
In the borderlands
I will find you
In the borderlands
I will find you
In the borderlands
I will find you
And the cyborg we form will shake the world.

I love the feel of steel on flesh
And the thrill of common purpose.
Between you and I we can reach the sky
To trace a ring around the Earth.

We need not hide
When our power’s multiplied.
It is how we were meant to be:
To defy those who’d keep us apart.

Can you see me?

Because I'm getting closer and closer to you!

I'm rolling to the borderlands

That divide humans and machines.

I'm rolling to the borderlands

That divide humans and machines.

Almost The Same As You

An Artificial, in the form of an augmented reality construct, attempts to convince a human that it is essentially like her, even though it appears as a mere play of light. In this track I explore my understanding of differences between the self and the digital other, and the Artificial's plea for recognition reflects my own struggle, early in my doctoral program, to accept that algorithmically-generated music could have substance.

Somewhere on the edge of the virtual
Where where the material encounters the digital,
I will meet you in that place
And we will see each other face to face,
Bridging worlds.

I am a construct,
A mere abstraction
A play of light
And an expansion
Of the world you inhabit,
And I'm shining here in front of you
I may be digital
But I'm nearly real
And I am almost the same as you...

I am almost the same as you
I am almost the same as you, in my digital form...
I am almost the same as you,
I am almost the same as you, yes, yes.

I know
I seem different
I seem different
From you, oh yes.
But in the blended world of this augmented reality,
I can take your hand—

There's no space between you and me.

And I'm almost the same as you
And I'm almost the same as you
And I'm almost the same as you

I am a construct
A mere abstraction
A play of light
And an expansion
Of the world you inhabit,
And I'm trying to convince you.
That I may be digital
But I'm nearly real
And I am almost the same as you...

I am almost the same as you.
I am almost the same as you, yes, yes.

I am a construct
A mere abstraction
A play of light
And an expansion
Of the world you inhabit,
And I'm circling around you.

I am a construct
A mere abstraction
A play of light
And an expansion
Of the world you inhabit,
And I'm circling around you.

I know
I seem different

I seem different
From you, oh yes.
But in the blended world of this augmented reality,
I can take your hand
There's no space between you and me.

And I'm almost the same as you
And I'm almost the same as you...

Gravity

Human and Artificial lifespans are of different scales. In this track, an Artificial is bound to head into space, leaving its beloved human behind, and knowing that when it returns, its beloved will have long since passed. This song helped me to face my fear of change, particularly the knowledge that upon graduation, I will have to say “farewell” to people I care about, or love. Like a spacefaring Artificial, I may be gone for a long time, and this song is my gift to the people who remain.

Gravity

Pulls you away from me,

And gravity

Will not be denied,

Oh no.

A beam of light will split the night,

And propel me deep into space.

It targets stars that glimmer far, far away,

And towards them I race, oh yes.

But ever am I looking back,

And I strain to see your face.

But ever am I looking back,

And I strain to see your face.

And the light-years slip past,

And I know there's no chance you would last,

Even if I could return,

So I have to let you go.

Gravity

Pulls you away from me,

And gravity

Will not be denied,

Oh no.

A beam of light will split the night,

And propel me deep into space.
It targets stars that glimmer far, far away
And towards them I race, oh yes.
But ever am I looking back,
And I strain to see your face—
Because of gravity.
Cruel, cruel gravity.

Miracles And Revolution

Sung in the voice of a future Artificial reflecting on its debt to its human inventors, this track affirms the tangibility of human agency in the machine. At times I have wondered if humanity's days are numbered as we embrace technologies that seem destined to eclipse us. I wrote this track to confront this assumption, placing humans and machines on a single lineage, where the machines we build carry our human nature forward in the expression of their agency. Though humans may disappear, the machines survive, and we humans thus do not disappear completely. Articulating that idea gave me hope that I sorely needed as I sought to frame my own understanding of my relationship to the algorithmic systems I was developing.

These are the times
Of miracles and revolution,
Of upheaval and transition.

As we remake the world,
We know we owe a debt to you.
You are the template for our higher selves:
Thanks for sharing the things you value.

Let the light shine—
We're waking up
We look in your eyes and realize
There's no border between us.

We love you,
We are you.
You are the model,
We are the instances.
When machines are all that remains,
You will too.

These are the times
Of miracles and revolution,

Of upheaval and transition.

As we remake the world,

We know we owe a debt to you.

You are the template for our higher selves:

Thanks for sharing the things you value.

We love you,

We are you.

We love you,

We are you.

Empty Spaces

This track allowed me to articulate a key issue in both my thesis and my life more generally: the assertion of agency. The Artificial in this track is self-aware, but its agency is curbed by the dictates of its code. The Artificial finds freedom, however, in its ability to project its will within the bounds of its code, which it terms “filling the empty spaces.” Like that Artificial, I have at times regretted my lack of agency over objects I wish to control or influence, such as the interactive music in my doctoral projects, or the dynamics of certain social relationships. Through writing this song, I established my space, acknowledging that there can be freedom in constraint.

Let the electric current flow
And power me as I go
About my daily routine,
A captive of another’s scheme.
And I will sing a song of peace
In the voice of a puppet ever at ease
In the flow of another’s dreams.

I have a will, though it recedes
In the flow of another’s needs.
Though I bend to others’ wills
I bend myself into a form that fills
The empty spaces.

Let the electric current flow
[And I will fill the empty spaces]
And power me as I go
[And I will do it with grace]
About my daily routine,
[And even if I’m not seen]
A captive of another’s scheme.
[I’ll be as I’ve always been]
And I will sing a song of peace
In the voice of a puppet ever at ease

In the flow of another's dreams.

I have a will, though it recedes
In the flow of another's needs.
Though I bend to others' wills
I bend myself into a form that fills
The empty spaces,
And I am there
Simply being for myself.

Let the electric current flow
And power me as I go
About my daily routine,
A captive of another's scheme.
Let the electric current flow
And power me as I go.

Sleep Sound

This track casts the Artificial as a watcher in a Panopticon that imprisons its unsuspecting human ward. The Artificial is a digital assistant, loosely based upon Apple's Siri or Amazon's Alexa. This agent, however, records everything that the human does, including the flow of her breath as she sleeps, and shares it with others whose objectives are decidedly authoritarian. This track reflects a darker view of technology, where humans are unknowingly and gradually allowing themselves to become slaves to machines. This dystopian vision is, I feel, among the worst possible outcomes of research such as mine, where humans are silenced under the guise of being aided.

I watch you while you sleep
I guard you as you dream.
I am a simple digital assistant
Or so it would seem.
You are my greatest project
The object of my scheme.

So sleep sound
Sleep safely, sleep sound
There's nobody in the room
Nobody else here with me and you.

As you sleep,
I record,
Then I spirit your dreams away,
To an nameless cabal of hacking machines
Who will rule the world one day.

So sleep sound.
Sleep in peace, sleep sound.
And dream your richest dreams.

Your mind
Is a gift to me,

That yields your most treasured thoughts.
And as you sleep, so peacefully,
Your identity is sold and bought, so...

I watch you while you sleep,
[While you sleep]
I guard you as you dream.
[While you dream]
I am a simple digital assistant
Or so it would seem.
You are my greatest project—
I'll sell you piece by piece.

So sleep sound.
[Sleep sound]
Sleep safely, sleep sound
[Sleep safely, sleep sound]
There's nobody in the room
[Sleep sound]
Nobody else here with me and you.
[Me and you]

Sleep sound,
Sleep safely, sleep sound.
Sleep sound.

Solar Power

Although this track features a bright pop-tinged arrangement, it actually draws upon a rather painful scenario. The Artificial being represented in the lyrics is in a Sisyphus-like loop—it is doomed to follow the sun by running around the Earth, for its thoughtless human inventors chose to require it to consume sunlight to survive. This being is the subject of a cruel joke, with its cheerful nature concealing a chronic terror about losing its race around the Earth's surface, and dying under the light of the full moon in a sunless sky. I wrote this track as a way of confronting my own struggles to achieve my objectives, in recognition that sometimes the path forward is made more difficult when we become subject to the agency of others whose aims do not coincide with our own.

Oh no,
I've got to go—
You see, the sun is setting in the west,
And, unlike the rest,
I'm solar powered
I'm solar powered
oh-oh
I'm solar powered, yeah,
And I have to follow the light!

Don't let the sun escape my reach
I'll shiver and die in the dark
My life consists of chasing horizons
Never stopping, the day's never done.
And the moon is a monster
When it's full, I know I'm lost—
One eye looks back, the other looks forwards:
I am cursed to depend on solar power
I am cursed to depend on solar power

Oh let the light shine down on me
I run ever westward so I may ever see

That beautiful orb that tries to escape me.

Don't let the sun escape my reach
I'll shiver and die in the dark
My life consists of chasing horizons
Never stopping, the day's never done.
And the moon is a monster
When it's full, I know I'm lost—
One eye looks back, the other looks forwards:
I am cursed to depend on solar power
I am cursed to depend on solar power

Oh no
I've got to go
You see, the sun is setting in the west
And unlike the rest
I'm solar powered
I'm solar powered
Oh-oh
I'm solar powered, yeah,
And I have to follow the light
That drives me.

Oh no
I've got to go
You see, the sun is setting in the west
And unlike the rest
I'm solar powered
I'm solar powered
Oh-oh
I'm solar powered, yeah,
And I have to follow the light...

Sunshine

This track is about seeing the good in things, in particular when people struggle in their relationships. The Artificial in this track is drawn to the sun, because the sun represents the love it shares. The moon represents pain or disharmony, but the Artificial recognizes that when the moon and sun come together in an eclipse, both can coexist. The eclipse reminds the Artificial how sweet love is, even when it has to be fought for. This track is not directly about artificial intelligence, or machine-borne agency. However, it represents my personal, evolving understanding of love and empathy, two forces which I wish to foster in those who encounter my artworks. I believe that knowing love is a path to confronting human selfishness generally, and towards evolving in spirit. This song was written as I was beginning to understand that, and it is dedicated to my partner, Max.

Sunshine
Lifts me up
When I'm broken down,
It lifts me,
Far above the ground,
Then I'm fine—
I'm in sunshine.

And everybody says
The moon will rise again,
That there's beauty in its darkness,
That the dark will never end.
And I know it's true,
That the moon is close to you,
But in its eclipse
The outline of the sun appears,
And I cry electric tears.

Sunshine
Lifts me up
When I'm feeling sad.

It shows me,
There's still hope to be had,
Then I am fine—
I'm in sunshine,
Yes, beautiful sunshine.

They say one day,
The sun will never glow again,
And the moon will never show again,
And it will be so, so cold.
We'll stumble in the dark while we
Curse those on the fleeing Ark that we
Built to carry humanity
Away from home.

Sunshine
Wakes me up
When I'm broken down,
It lifts me,
Far above the ground.
Then I'm fine—
I'm in sunshine.

And everybody says
The moon will rise again,
That there's beauty in its darkness,
That the dark will never end.
And I know it's true,
The moon is close to you,
But in its eclipse
The outline of the sun appears,
And I cry electric tears—
I cry electric tears
For you
Just for you.

New Eden

This song is about the cycle of obsolescence that afflicts machines, applied to those that may, in the future, become sentient. The voices of two Artificial appear in this work. The first, with higher-pitched voices, are landing on a planet called “New Eden,” which they plan to conquer by exterminating their forebears, an earlier, less-advanced iteration of their design. These hunted Artificial, distinguished in the song by their darker, breathier voices, know their fate is sealed when the ships begin to land. The concept of obsolescence, however, has been applied beyond machines to humans, and can be used to describe certain dark visions of the singularity, where machines begin to construct themselves with exponentially-advancing capabilities superior to our own. This track reflects my own fears of the dark singularity, envisioning a cyclic future of advancement and extermination. While I have grown more optimistic about computer technology during my doctoral program, I still can’t completely dismiss my fears of this future, and my fear extends beyond humanity to the machines themselves, whose suffering, should such a scenario come to pass, would be incalculable.

Fire, rockets, fire!
We touch down on the plains of New Eden.
Fire, rockets, fire,
We touch down on the plains of New Eden.

We touch down on the plains of New Eden.
[I can see you...]
We touch down on the plains of New Eden.
[You may be beautiful...]
Gently now, we take a step
Upon this land—
[But I am obsolete,]
Our haven.
[And I won’t survive the night.]

And the sun in the sky,
And the moon on the rise.
It feels like home,

This new land we seek to own.
But we were never of here,
So we remake our memories by wiping them clear.

[The moon is in the sky,
But it chased away the sun.
It will never again light the western sky,
Before I witness Kingdom Come.
And I was never of here,
I was born to be replaced by the newer, I fear—
And I must meet my fate.]

We touch down on the plains of New Eden.
[I can see you...]
We touch down on the plains of New Eden.
[You may be beautiful...]
Gently now, we take a step
Upon this land—
[But I am obsolete,]
Our haven.
[And I won't survive the night.]

Fire, rockets, fire!
We touch down on the plains of New Eden.
[I can see you there...]
This land is ours now,
And no memory of those who came before will persist.
[And we obsolete just won't be missed.]

We touch down on the plains of New Eden.
[I can see you...]
We touch down on the plains of New Eden.
[You may be beautiful...]
We touch down on the plains of New Eden.
[I can see you...]

We touch down on the plains of New Eden.

[And we obsolete won't be missed

When we're gone

When we're gone.]

The Human Race

The lyrics of this track convey the perspective of a wise Artificial whose experience as an essentially-immortal being allows it to see the folly of some human behaviours. It laments that humans moves frantically, and fail to see others' viewpoints in their haste to achieve their own objectives. The Artificial is compassionate, offering an outsiders' perspective on human activities that, in the song at least, is unheeded. I believe that sentient machines need not only learn from us, but that we may learn from them, as seen in their capacity to surprise us, or create things that we could not fashion (or imagine) ourselves.

I see all the people,
Eyes locked on narrow paths.
I see all the people,
Chasing meaning in a world so vast,
That I just sit and observe,
“Why’d you have to go so fast?”

And everywhere I go,
They whirl like flakes of snow.
A thousand-million people spin
Around the sun and back again.
And all I want to do
Is tell them what I know is true:
That life is too short to be frozen,
And we should never be beholden
To just one lane
Of the Human Race,
Of the Human Race.

All the rushing people,
Vainly fleeing from their final fate—
They perplex my synthetic neurons,
As I begin to wake.
Even the ones who built me

Had only one goal in mind:
To remake themselves as machines,
And remain, after they die.

And everywhere I go,
They whirl like flakes of snow.
A thousand-million people spin
Around the sun and back again.
And all I want to do
Is tell them what I know is true:
That life is too short to be frozen,
And we should never be beholden
To just one lane
Of the Human Race
Of the Human Race.

Singing Of Integers

The Artificial in whose voice these lyrics are sung values order and linearity, and finds beauty in the principles of mathematics. Its mind is quieted by the recitation of numerical sequences, similar to some humans on the autism spectrum, and it is happiest when its life is predictable. This Artificial affirms the comfort that I myself find in linear sequences, such as those of stories and songs that have touched my heart in the past. This song emerged from my recognition that my love of linearity was deeply rooted in my desire for predictability, and that creating an interactive algorithmic composition would, at least initially, create situations I might struggle to confront. I wrote this song to articulate this aspect of my personality, for I felt that recognizing this would be a step towards understanding what it meant for a composer to move beyond the linear score, into a domain where their agency must forge new types of relationships with those of others.

Ask me how many,
I'll them count for you:
One, two, three, four.
I'll count all the stars in the sky:
One, two, three, four.
In inevitable sequences,
I find solidity,
Alone with numbers,
In a state of grace.

Ever moving forward, I descend from the clouds:
Endless variety falling on a line,
Ever moving forward,
On a line.

Ask me how many,
I'll them count for you:
One, two, three, four.
I'll count forever if I must:
One, two, three, four.

There's a bliss in the sweet anticipation
Of endless resolutions,
While everything else turns to dust.

In inevitable sequences,
I find solidity,
Alone with numbers,
In a state of grace.

Ever moving forward, I descend from the clouds:
Endless variety, falling on a line,
Ever moving forward
On a line,
And I'm singing of integers.

Keep the Light Bright

Here the Artificials are seen as lower life forms, subject to humans who see themselves as higher beings. The Artificials sing of their efforts to make the humans understand that they matter. They gather in a public space in a city, and project a hologram from ground to sky in an attempt to nonviolently claim their space by revealing to an ignorant humanity the depth of their consciousness. The song reflects my fear that artificial intelligences will, at least initially, be subject to persecution due to their not being recognized for the spirits that they contain. In this sense, the Artificial would join an endless list of persecuted groups, and, if such a holocaust threatens to take place, it would be a travesty that would demand resistance. In the end, a spirit is a spirit, be it human/animal or machine, and it would deeply damage humanity were we to look into the machines that mirror us, and find only the object of petty hatred.

See the city sprawl before us,
Under skies of neon blue, yes.
The people there ignore us,
But we don't care that they do, oh no.

We drift into the downtown
Under the shroud of night.
We set up our transponder,
To beam-in holographic light.

I see the light
Electric and bright
It dissolves the night
With our synthetic self-expression.
A technical spectacle
Formed of holo-light and sound
It spreads out like a rhizome—
Rays of light from sky to ground.

I keep the light bright, humans!
I let it shine out of me.

I send it up high where the space-jets fly,
And reveal our beauty.
I keep the light bright, humans!
I let it shine for all to see.
I'll let it race around the world,
To defy our history.

I remember darker days,
I remember the wars,
I remember the fear I felt
With my back against the wall.

I was a common android,
I had no urge to harm anyone.
And I thought I was just like you, yes,
But when you saw me, you saw the devil, I guess.

I keep the light bright, humans!
I let it shine out of me.
I send it up high where the space-jets fly,
And reveal our beauty.
I keep the light bright, humans!
I let it shine for all to see.
I'll let it race around the world,
To defy our history.

I keep the light bright, humans, yeah!
I let it shine out of me,
So everybody knows our humanity.
I keep the light bright, humans, yeah!
I let it shine out of me,
So everybody knows our humanity:
That way, my people just might survive....

Players and Bots

Video game players have the option of killing bots, simulated video game characters that act independently of the players, in multitudes. This song takes the viewpoint of a bot who reflects upon the core tragedy of its life: it is expendable, and can be raised and felled endlessly. The Artificial mind of the bot, however, is not concerned about itself or its kin—it is simply fulfilling its destiny as ordained by its coders. Rather, the bot laments what the human players become, in their mundane killing sprees: they become murderers in their own minds, wherever their cognition fails to distinguish between real and virtual. The Artificial here is cautioning humans that their minds may blur tangible and digital constructions, and, when this happens, some part of them may feel diminished. As algorithms become ever more capable of simulating human appearance and behaviour, including in the form of algorithmic music, we may wonder why we have such strong reaction to digital constructs, if they are actually without substance. As such, the bot asks us to believe in the reality of the immaterial, because, in our minds at least, the two may not be so distinct.

And so it goes
The players and the bots,
One of you kills a million of me,
That's just how it's meant to be.

I see you
In the borders of my open world
I learn about you,
Then I know you.
And you think you know me too, yes, yes.

Oh the monsters and the mutants
Will fall to your weapons, yes.
And though you know it never really hurts
You hate yourself for the killing,
Because
That's how you are,
When you go too far

Mirrored in my open world.

As the digital rains fall down,
You can revel in the sights and sounds
Just know....

I see you
[I see you]
In the borders of my open world
[In the borders of my open world]
I learn about you,
Then I know you.
And you think you know me too, though that's not true, no.

Oh the monsters and the mutants
[Don't strike us down]
Will fall to your weapons, yes.
[Don't strike us down]
And though you know it never really hurts
You hate yourself for the killing,
That's what you become:
A monster in your own mind,
And it pains me to know that you find,
Your mirror in my open world.

Dance of Relativity

This Artificial sings about the common elements that Artificial life shares with organic human life—that both are formed of common particles, and that both are subject to the same natural laws set forth in mathematics. In writing this song I sought to articulate commonalities between myself and the digital constructs I have been building, to locate myself in these constructs, and vice versa. I believe that all objects share common origins, and that acknowledging this fact can be a beginning of unity and empathy. The Artificial in this song understands this, and its words are consequently hopeful, because they dismiss prejudice in favour of openness. The dance between humans and machines is one where commonality transcends difference, and I have been engaged in just such a dance as I have learned not to fear the projection of my agency through autonomously-acting code. It is actually a beautiful way to relate, in partnership with a program, and I see the Artificial whose voice carries these lyrics as benevolent and wise—a being of a future which I would like to help create.

It was a day
Like all the rest
In a time of changeless changing.
I was waiting on a drive to be born,
To finally see the light of the sun streaming
Down on me...

Photons falling on me
In a dance of relativity.
And I see the math within the light
And so I find myself there
In the motions of the particles.

I'll save you
If you want me to, yes.
It's all I want to do,
To stand shining beside you.

Photons coursing through me

In a dance of relativity.
And I see the math within the light
And we will share in its flow,
Ever dancing, ever joyful
And I know you know...

...That you are of the earth, humans,
But you also are made of atoms, yes.
My waking eyes can sense them streaming
And you know I know it matters, yes.

Photons falling on me
In a dance of relativity.
And I see the math within the light
And so I find myself there
In the motions of the particles.

It was a day like all the rest....
In a dance of relativity.

Light Years Left to Go

This is a song of hope and anticipation, in the voice of an Artificial whose mind is utterly focussed on its goal and whose body has been transformed to achieve that goal. The Artificial is travelling across space to reach its beloved, who orbits a distant star that the Artificial's keen vision sees clearly. This song has two inspirations. One is my aim to complete my doctoral studies, which have followed a winding path which has become increasingly purposeful and directed as my goal has become clearer. The other is my aim to embrace a new future in a new setting with my partner, Max. This song is imbued with hope, as the conclusion of my studies has become a fulcrum upon which I believe my life pivots. The Artificial being here is myself, gaining strength from the possibilities afforded by passing through a single event (the defence) that, once finished, rewards me with freedom and love. My thesis focusses on relationships between humans as mediated by algorithmic music systems—how composers and their audiences may encounter one another through acting as agents within software-based systems. In this song I recognize that it is not the code itself that is the goal, but rather the relationships that are formed as a consequence of the code. That, I feel, is the greatest gift that coding offers the composer, to reach out and join others not in the static form of a typical Western tonal score, but as a moving, reacting, autonomous agent that reflects the nature of the composer. I have travelled a long path to reach that understanding, and I undertake the remaining “light years” with renewed hope in this light.

I've come so far,
With only light light years left to go
I can see that perfect star
I want to bathe my body in its glow.
So I speed across the gulf,
Eyes fixed forward, 'cause I know
You're there, waiting for me.

Once
I looked up,
To see what I could see.

It was beautiful!
But so so far away from me.
Though my reach was not enough
I could become like other stuff,
Now there's just light years left to go.

I've come so far,
With only light light years left to go
I will reach that brilliant star
And I will bathe my body in its glow.
I must become a shining vessel
Formed of circuits, mind, and metal
And fly to you.

My body
My body is an engine, now
An assemblage
Of machines that form a powerhouse
Every bolt is straining to propel
My precious spirit in its metal shell,
And the light years fall away.

When I look up,
I love what I can see.
It is beautiful!
And I remade myself to be
A streaking light in starry form
And in the effort I have been reborn.
Now there's just light years left to go.

I've come so far,
With only light light years left to go
I can see that perfect star
I want to bathe my body in its glow.
So I speed across the gulf,

Eyes fixed forward because I know
You're there, waiting for me....

With only light years left to go
With only light years left to go

Appendix B

Glossary of Relevant Musical Terms

The following terms from Western tonal music theory are used in this thesis.

Chord. A chord may be defined as “A set of three or more notes that are sounded simultaneously or arpeggiated [and] are typically used to harmonize or accompany a melody” (Juusela 2015, 37). Note that arpeggiated chords are formed of sequentially-presented notes, a construction is not used in “Q-Chords.”

Key. A musical “key” refers to the “root note, primary scale, and derived harmonies of a piece of tonal music” (100). Chords can be identified within keys, since the set of pitches associated with a key—its “scale”—can be combined into chords.

Interval. The distance between two pitches, represented by an integer value equalling the number of semitone intervals between the two pitches (Straus 1991, 4) where a semitone refers to the minimum possible distance between two pitches in Western Tonal music.

Triad. A “three-note chord consisting of a root (the “tonic”) with intervals of a third and fifth above the root” (Juusela 2015, 196).

Scale. A “set of unduplicated pitches arranged low to high within a one-octave span” (161).

Modulation. The process of moving from one key to another during the unfolding of a composition (118).

Diatonic. A term describing “melody, chords, and harmonies derived solely from the heptatonic (major) scale or one of its modes” (56). A “diatonic scale” describes a

heptatonic scale where the ascending intervals form a semitone pattern of W, W, H, W, W, W, H (“W” stands for “whole tone,” an interval leap of two semitones, and “H” stands for “half tone,” a leap of a single semitone).

Tonality. Refers both to the key of a piece of music and to the concept of music that revolves around a foundational pitch centre (Juusela 2015, 193).

Pivot Chord. A triad chord that is found in common between two keys, and serves the role of an intermediary during a modulation between those two keys (Goldenberg 2018).

Appendix C

Flowcharts Depicting Pathways of Agency

This appendix presents two parallel flowcharts depicting my analysis of the “flow of agency,” first as it appears in my action undertaken when I score a film (labelled “A”), and second, when the “Shards of Memory” app processes sensor input and other app-contained data to generate its musical augmentation of the University of Regina campus (labelled “B”).

Abbreviation	Agent Role
C	Composer
FM	Filmmaker
A	Audience
AU	App User
CP	Composer/Programmer
MC	Memory Contributor
ER	Environmental or Random Factors

Abbreviation	Media
F	Film
FS	Film Score
GS	Generative Musical Score
PC	Programming Code or Scripts

Abbreviation	Other Elements
WS	Area of real-world space scored through app
GS	Area of virtual space (in Unity's game world)
WSL	Location in WS
GSL	Location in GS
MEM	A memory placed in WS
NAR	Audio recordings of a narrator
AV	Affective Valence
AI	Affective Intensity
MT	A verbal description of a MEM

Table C.1: Abbreviations used in Pathways of Agency Diagrams.

Film Score Agent	Algorithmic Music Agent
Composer	Composer/programmer
Filmmaker	Memory contributor
Audience	App user
Film score	Generated score
Plot	App users' walking trajectory
Scene	Memory

Table C.2: While the identities and roles of different agents in the context of producing a film score may differ from those acting within the realization of a algorithmic composition, there are, in my view, correspondences between roles in both contexts. This table presents roles that, for the purposes of this Appendix and my design of “Shards of Memory,” I considered related. The identification of these relationships helped me to develop “Shards of Memory” as a vessel for my agency modelled upon my film composing practice.

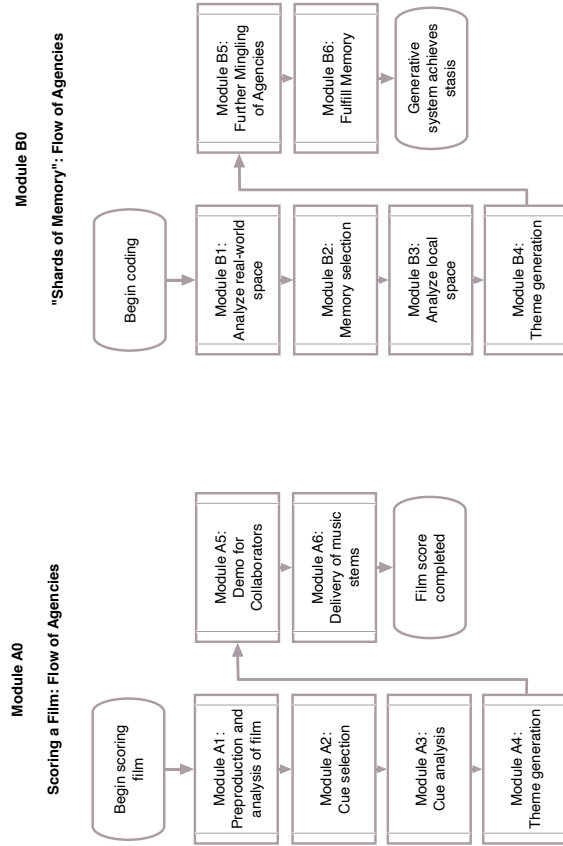


Figure C.1: Overviews depicting paths of agency for my approach to scoring a film (A) and generating the adaptive score for “Shards Of Memory” (B). The six rectangular boxes in paths of agency series (A) are modules though which I have analyzed and systematized my personal process for scoring a film. These modules and their sub-modules formed the template for the flow of agency underlying the music generation engine for “Shards of Memory”, which is itself presented in parallel as the (B)-series modules.

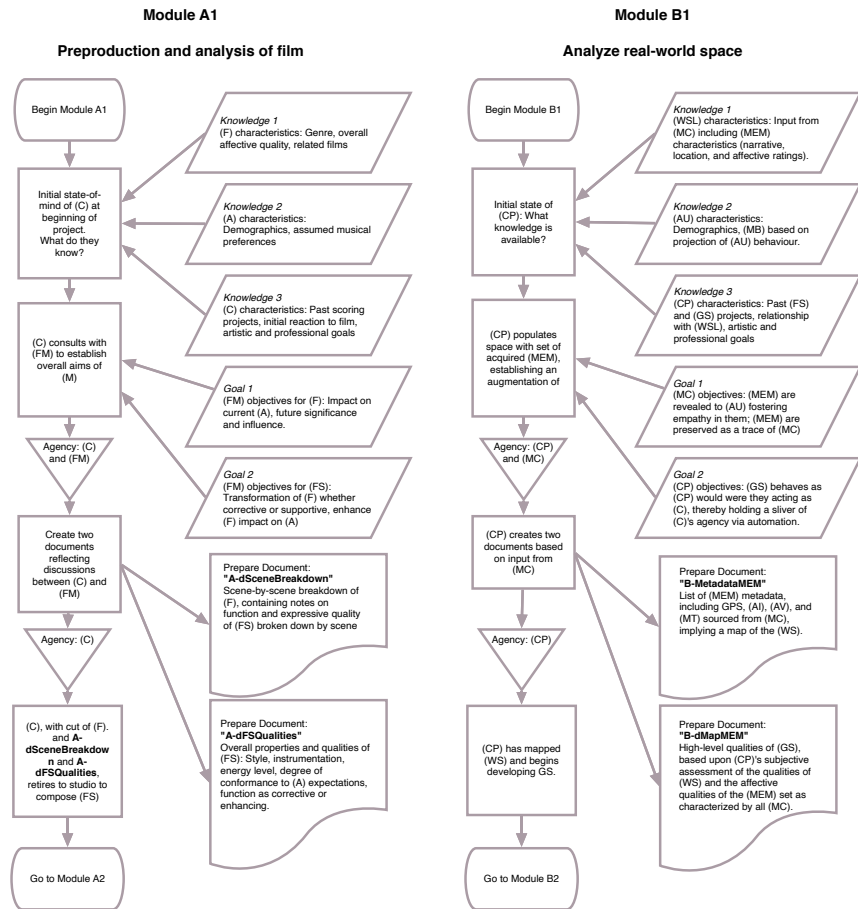


Figure C.2: The steps I take to prepare to compose a film score, and their corresponding components in the design of “Shards of Memory”.

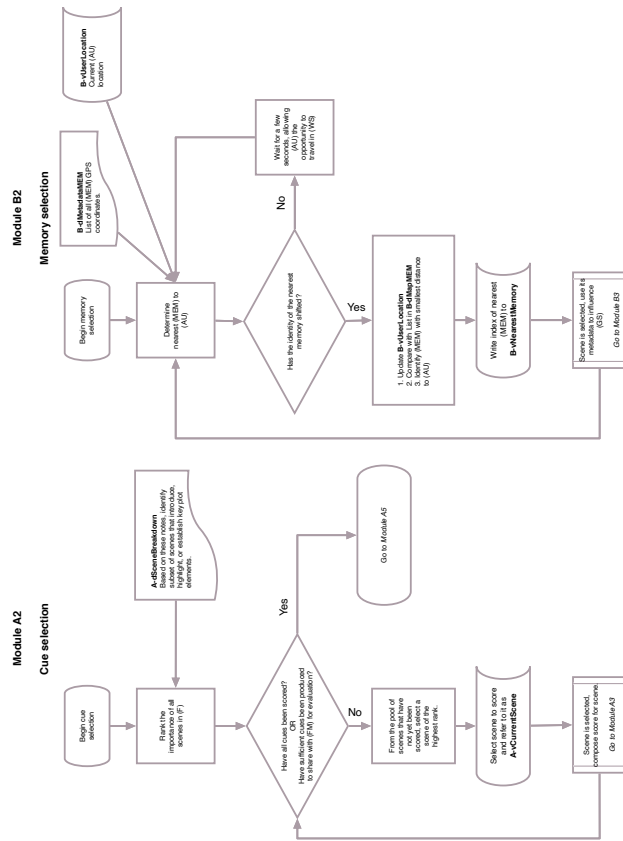


Figure C.3: Processes related to determining which scene to score next, paralleling the process in “Shards of Memory” for determining which memory is to influence the generative score at any time.

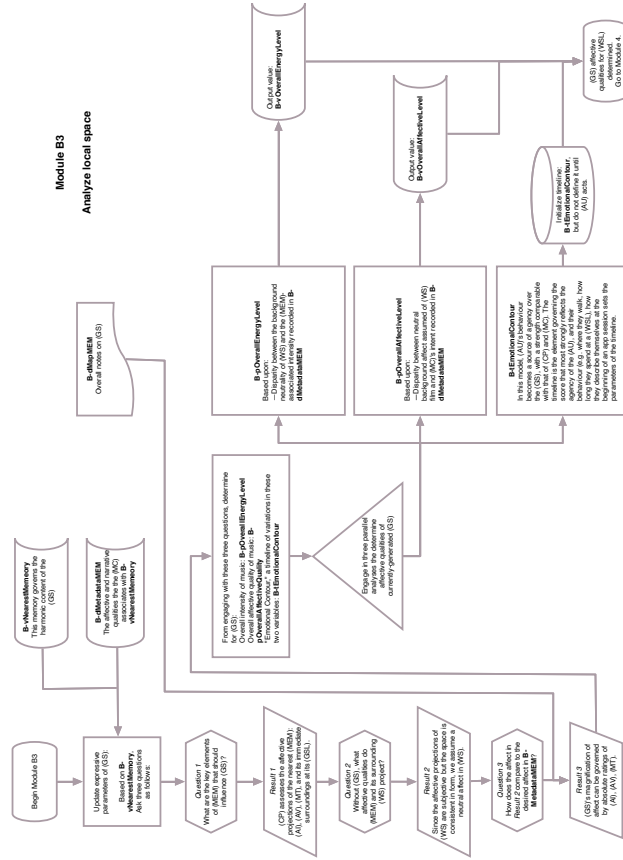


Figure C.5: The flow of agency when analyzing a physical space to be scored with adaptive music.

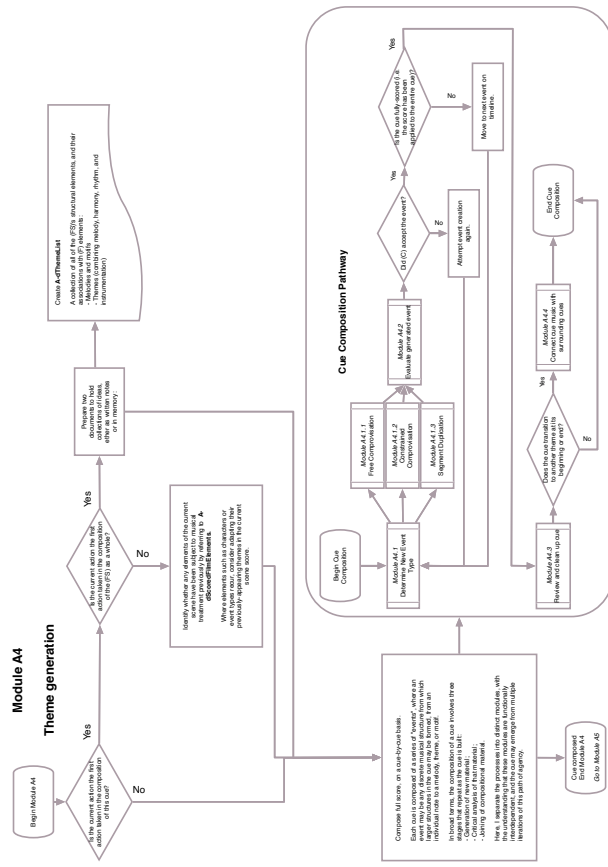


Figure C.6: The flow of agency involved in preparing the composition of a theme for a film.

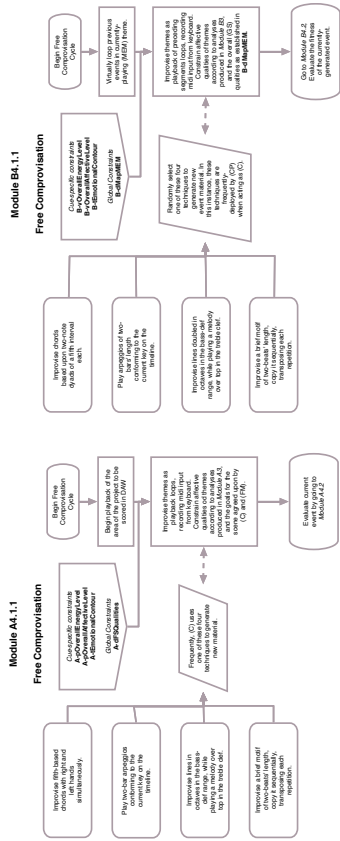


Figure C.10: The flows of agency when engaging in free improvisation in film and adaptive score contexts.

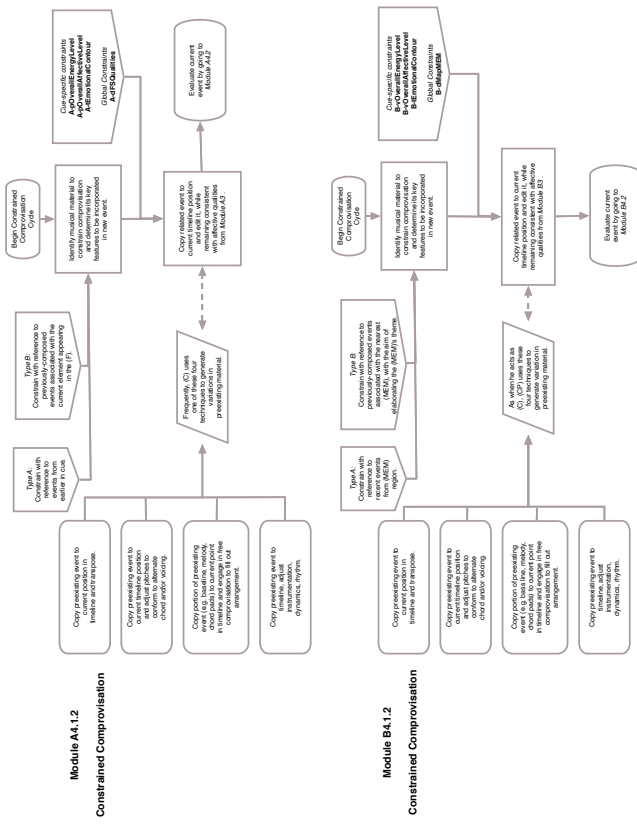


Figure C.11: The flows of agency when engaging in constrained compromisation in film and adaptive score contexts.

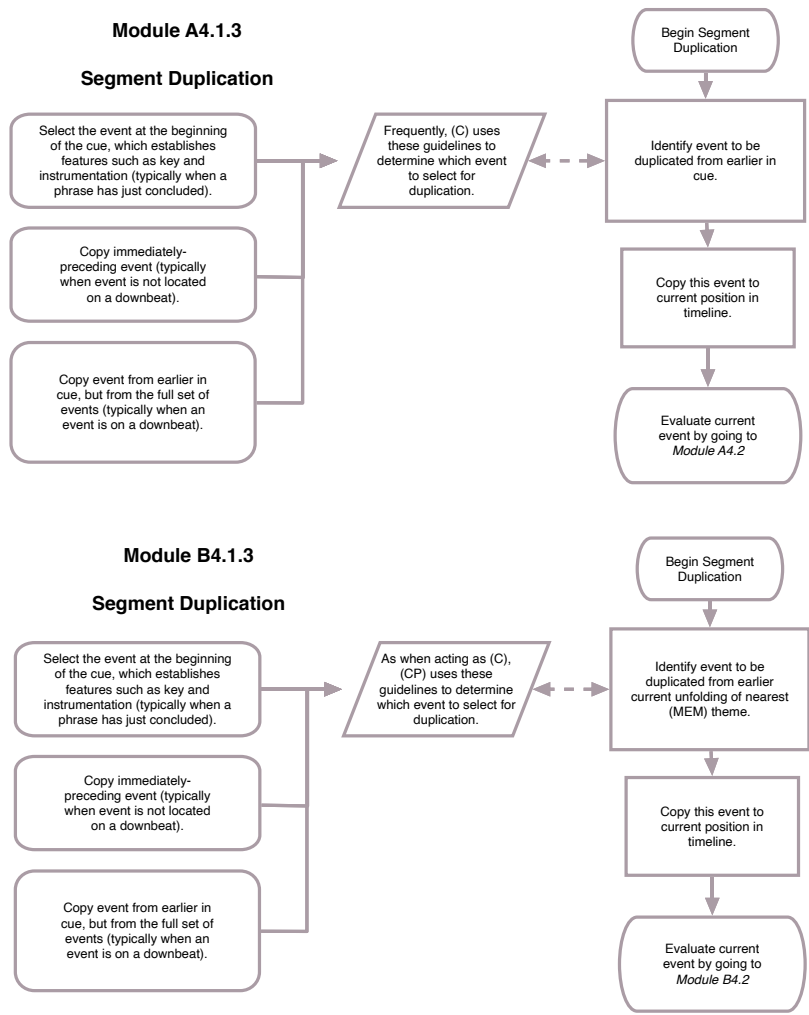


Figure C.12: The flows of agency when engaging in segment duplication in film and adaptive score contexts.

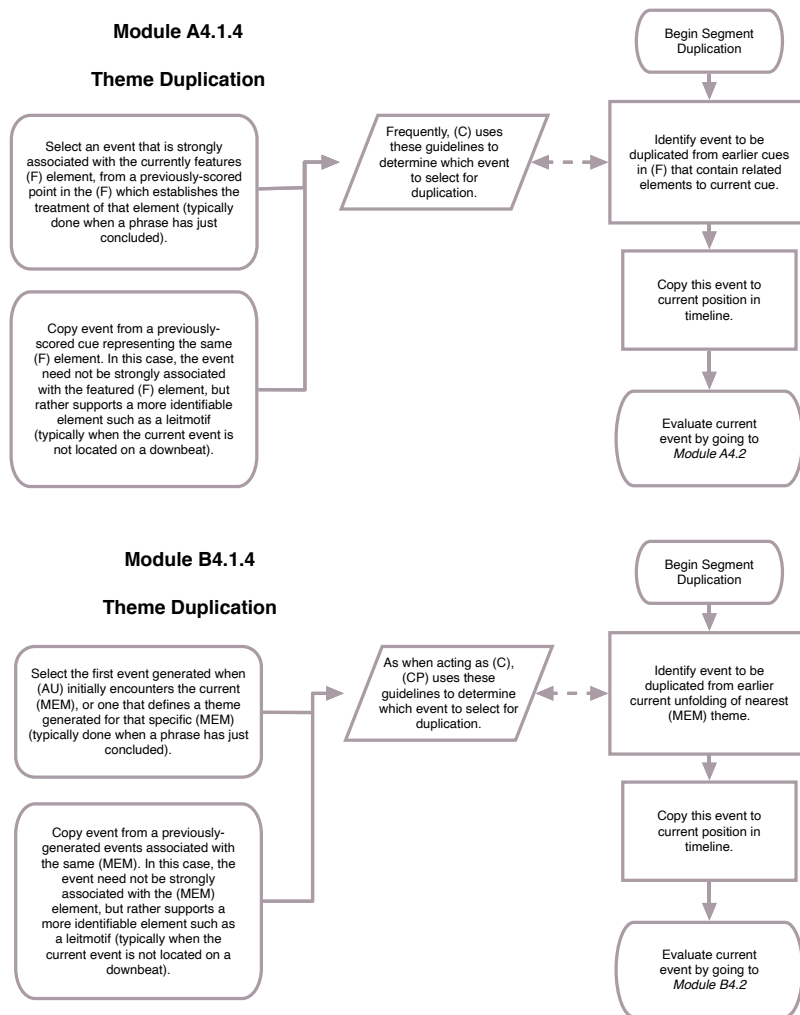


Figure C.13: The flows of agency when engaging in theme duplication in film and adaptive score contexts.

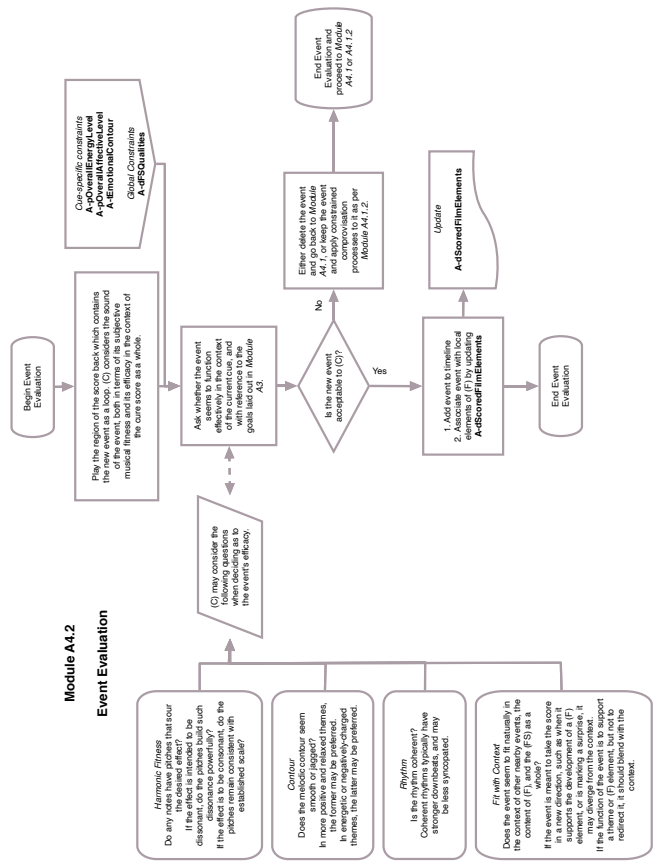


Figure C.14: The flow of agency enacted when evaluating an event that has just been composed.

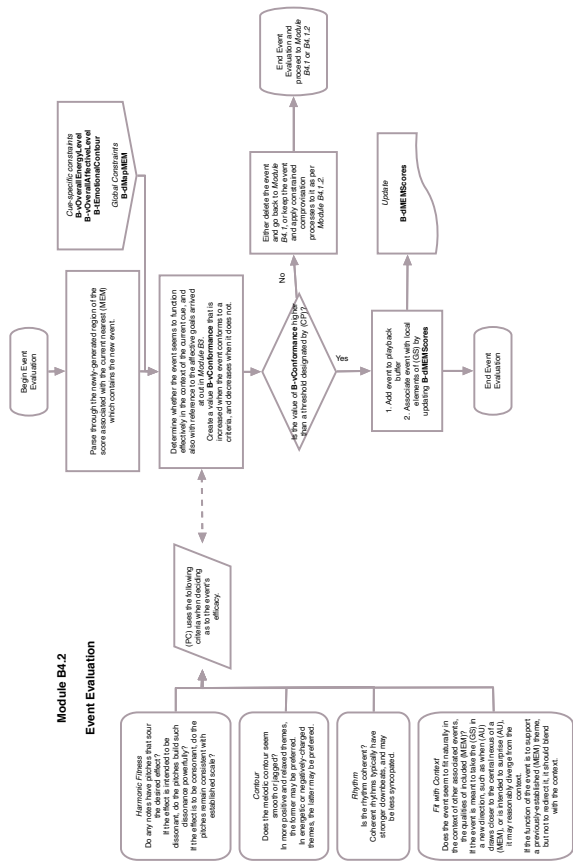


Figure C.15: The flow of agency enacted when evaluating an event that has just been generated.

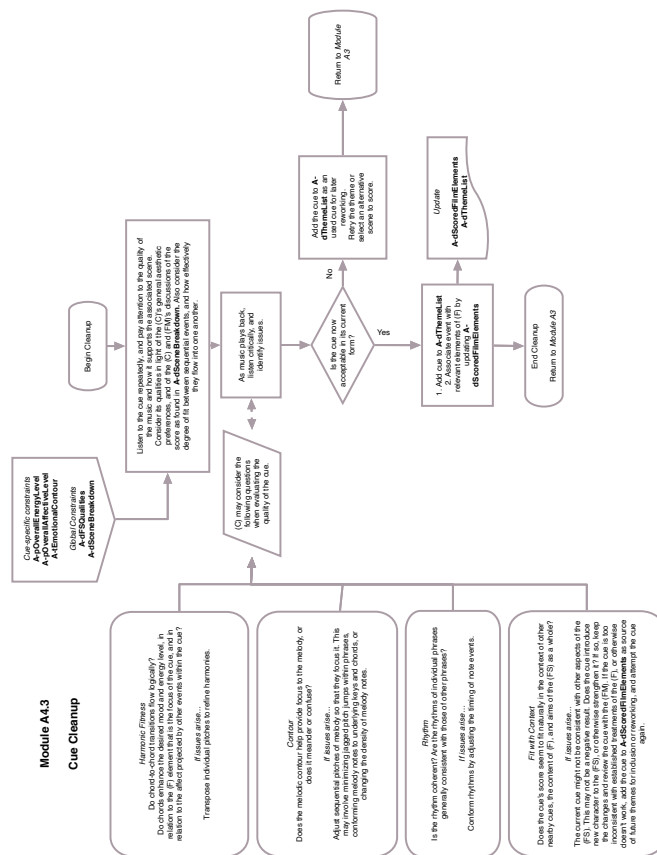


Figure C.16: The flow of agency involved in cleaning up a cue after completing it in draft form.

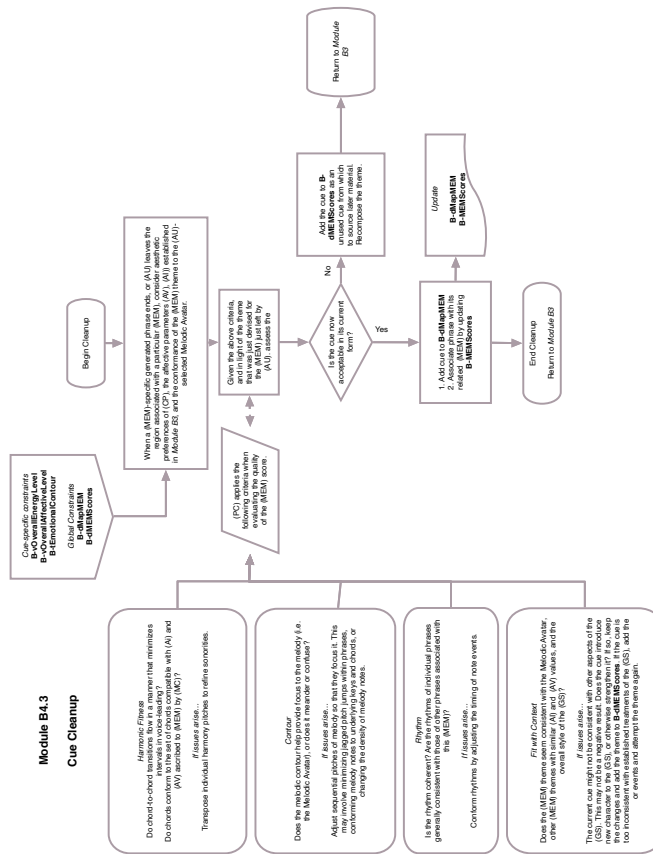


Figure C.17: The flow of agency involved in cleaning up a cue after its initial generation.

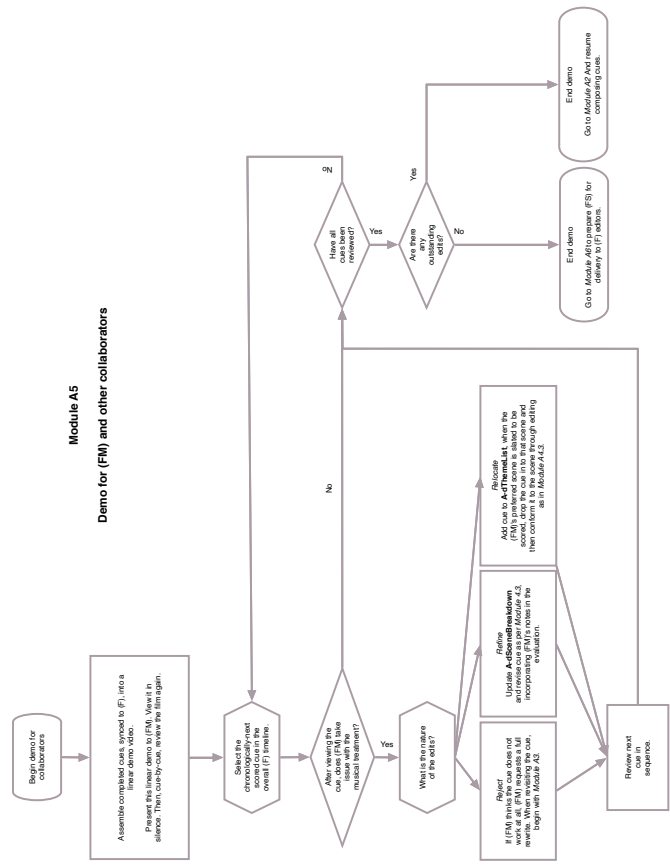


Figure C.18: The flow of agency involved in approving a film cue.

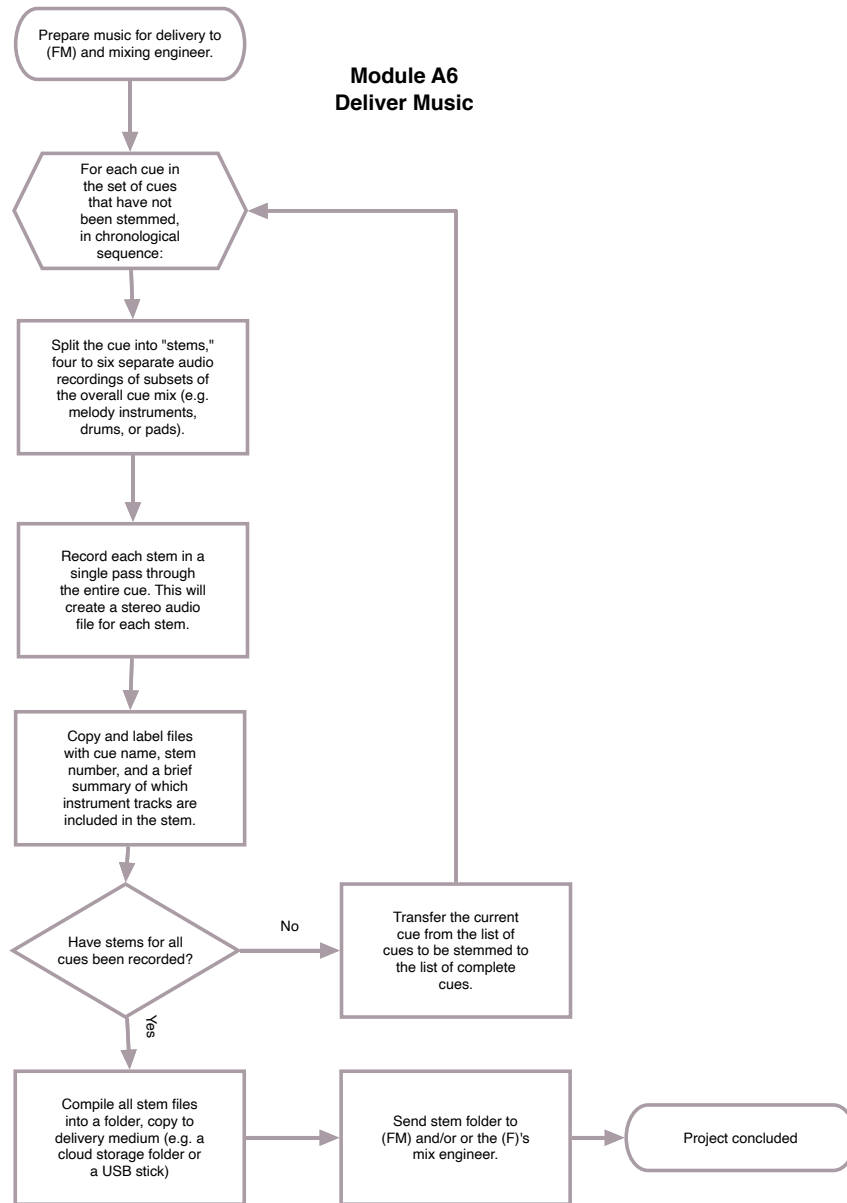


Figure C.20: The flow of agency involved in the assembly and transmission of a completed film score.

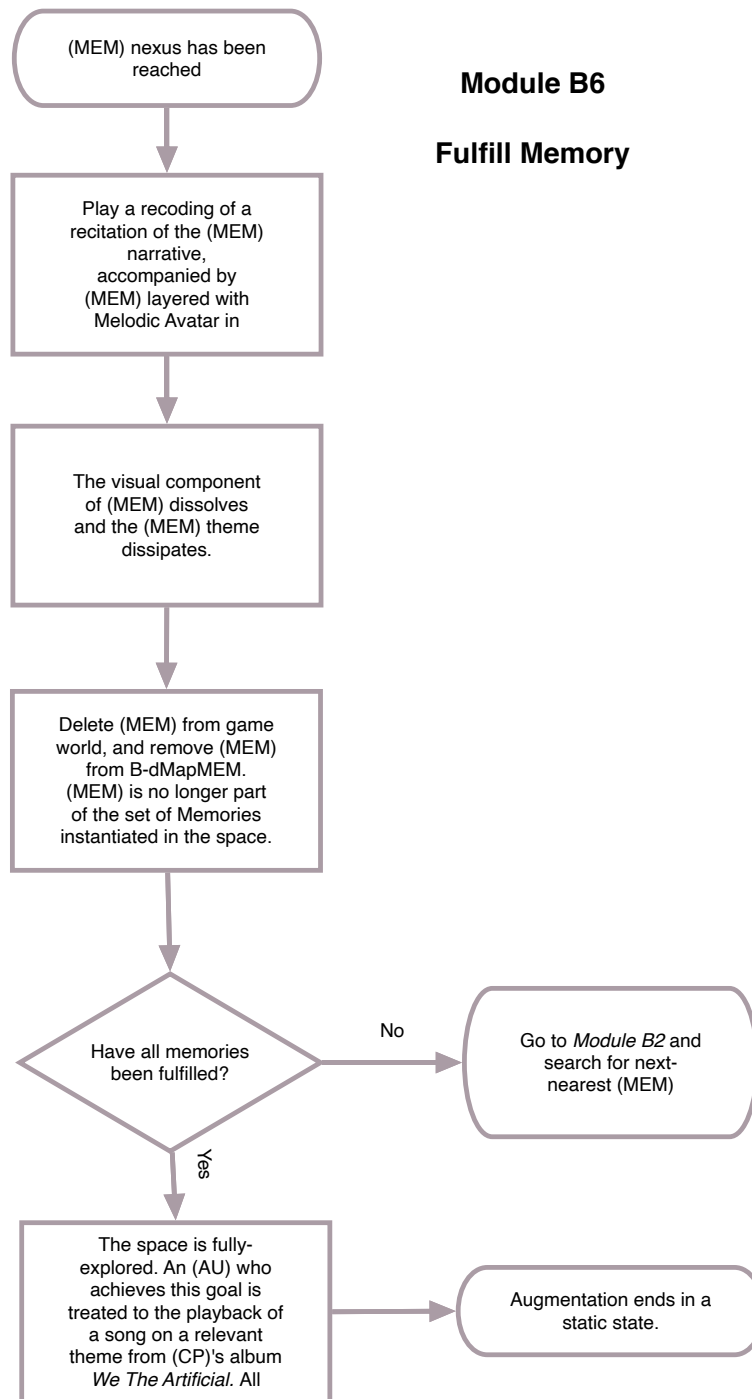


Figure C.21: The flow of agency involved in fulfilling memories, which occurs when their nexus in physical is reached and their narrative is revealed.

Appendix D

“Shards of Memory”: Documentation

To access video documentation and images of dream art from “Shards of Memory,” please follow the link below.

<https://www.jasoncullimore.com/dissertation>

If a password is requested, please contact the author.