

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

**L'ARTE DI INTERAZIONE MUSICALE:
NEW MUSICAL POSSIBILITIES THROUGH
MULTIMODAL TECHNIQUES**

JORDAN NATAN HOCHENBAUM

A DISSERTATION
SUBMITTED TO THE
VICTORIA UNIVERSITY OF WELLINGTON AND
MASSEY UNIVERSITY IN FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN SONIC ARTS

NEW ZEALAND SCHOOL OF MUSIC

2013

Supervisory Committee

Dr. Ajay Kapur (New Zealand School of Music)

Supervisor

Dr. Dugal McKinnon (New Zealand School of Music)

Supervisor

© JORDAN N. HOCHENBAUM, 2013

NEW ZEALAND SCHOOL OF MUSIC

Abstract

Multimodal communication is an essential aspect of human perception, facilitating the ability to reason, deduce, and understand meaning. Utilizing multimodal senses, humans are able to relate to the world in many different contexts. This dissertation looks at surrounding issues of multimodal communication as it pertains to human-computer interaction. If humans rely on multimodality to interact with the world, how can multimodality benefit the ways in which humans interface with computers? Can multimodality be used to help the machine understand more about the person operating it and what associations derive from this type of communication?

This research places multimodality within the domain of musical performance, a creative field rich with nuanced physical and emotive aspects. This dissertation asks, what kinds of new sonic collaborations between musicians and computers are possible through the use of multimodal techniques? Are there specific performance areas where multimodal analysis and machine learning can benefit training musicians? In similar ways can multimodal interaction or analysis support new forms of creative processes?

Applying multimodal techniques to music-computer interaction is a burgeoning effort. As such the scope of the research is to lay a foundation of multimodal techniques for the future. In doing so the first work presented is a software system for capturing synchronous multimodal data streams from nearly any musical instrument, interface, or sensor system.

This dissertation also presents a variety of multimodal analysis scenarios for machine learning. This includes automatic performer recognition for both string and drum instrument players, to demonstrate the significance of multimodal musical analysis. Training the computer to recognize who is playing an instrument suggests important information is contained not only within the acoustic output of a performance, but also in the physical domain. Machine learning is also used to perform automatic drum-stroke identification; training the computer to recognize which hand a drummer uses to strike a drum. There are many applications for drum-stroke identification including more detailed

automatic transcription, interactive training (e.g. computer-assisted rudiment practice), and enabling efficient analysis of drum performance for metrics tracking.

Furthermore, this research also presents the use of multimodal techniques in the context of everyday practice. A practicing musician played a sensor-augmented instrument and recorded his practice over an extended period of time, realizing a corpus of metrics and visualizations from his performance. Additional multimodal metrics are discussed in the research, and demonstrate new types of performance statistics obtainable from a multimodal approach.

The primary contributions of this work include (1) a new software tool enabling musicians, researchers, and educators to easily capture multimodal information from nearly any musical instrument or sensor system; (2) investigating multimodal machine learning for automatic performer recognition of both string players and percussionists; (3) multimodal machine learning for automatic drum-stroke identification; (4a) applying multimodal techniques to musical pedagogy and training scenarios; (4b) investigating novel multimodal metrics; (5) lastly this research investigates the possibilities, affordances, and design considerations of multimodal musicianship both in the acoustic domain, as well as in other musical interface scenarios. This work provides a foundation from which engaging musical-computer interactions can occur in the future, benefitting from the unique nuances of multimodal techniques.

Contents

Chapter 1 Introduction.....	1
1.1 NOISE AND INSPIRATION.....	1
1.2 ON HUMAN INTERACTION.....	2
1.3 A DEFINITION OF MULTIMODALITY	4
1.4 OVERVIEW	8
1.5 SUMMARY OF CONTRIBUTIONS	10
Related Work.....	12
Chapter 2 Background and Motivation	13
2.1 A BRIEF HISTORY OF MULTIMODALITY AND HCI.....	13
2.1.1 Detecting Affective States	15
2.1.2 Selected Examples of Multimodal Musical Systems.....	16
2.1.3 Summary.....	19
2.2 A HISTORY OF RELATED PHYSICAL COMPUTING.....	19
2.2.1 New Interfaces & Controllers: Building On and Diverging From Existing Metaphors	20
2.2.2 Hyperinstruments	23
2.3 TOWARDS MACHINE MUSICIANSHIP.....	26
2.3.1 Rhythm Detection	26
2.3.2 Pitch Detection	28
2.4 MUSIC AND MACHINE LEARNING	29
2.4.1 Supervised Learning and Modeling Complex Relationships in Music	30
2.5 SUMMARY.....	32
Research and Implementation.....	34
Chapter 3 The Toolbox.....	35
3.1 INSTRUMENTS, INTERFACES, AND SENSOR SYSTEMS	35
3.1.1 Esitar.....	36
3.1.2 Ezither	37
3.1.3 Esuling.....	38
3.1.4 XXL	39
3.2 NUANCE: A SOFTWARE TOOL FOR CAPTURING SYNCHRONOUS DATA STREAMS FROM MULTIMODAL MUSICAL SYSTEMS	40

3.2.1 Introduction to Nuance.....	40
3.2.2 Background and Motivation	42
3.2.3 Architecture and Implementation.....	44
3.2.4 Workflow	50
3.2.5 Summary	51
Chapter 4 Performer Recognition	53
4.1 BACKGROUND AND MOTIVATION	53
4.2 PROCESS	55
4.3 SITAR PERFORMER RECOGNITION.....	55
4.3.1 Data Collection	56
4.3.2 Feature Extraction.....	58
4.3.3 Windowing.....	59
4.3.4 Classification.....	59
4.3.5 Results and Discussion	60
4.4 DRUM PERFORMER RECOGNITION.....	64
4.4.1 Data Collection	65
4.4.2 Feature Extraction.....	67
4.4.3 Understanding Data Through Multimodal Visual Feature Clustering.....	68
4.4.4 Classification.....	70
4.4.5 Results and Discussion	71
4.5 DISCUSSION.....	76
Chapter 5 Drum-Stroke Computing.....	79
5.1 BACKGROUND AND MOTIVATION	79
5.2 DATA COLLECTION	82
5.3 ANALYSIS FRAMEWORK	82
5.3.1 Surrogate Data Training	83
5.3.2 Onset Detection.....	83
5.3.3 Feature Extraction.....	85
5.4 DRUM HAND RECOGNITION	86
5.4.1 Classification.....	86
5.4.2 Results: About the Tests.....	86
5.4.3 Results: Test One – All Data (Individual vs. Combined Scores)	86
5.4.4 Results: Test Two – Data Split.....	89

5.4.5 Results: Test Three – Leave One (performer) Out	90
5.5 DRUM PERFORMANCE METRICS	91
5.5.1 Cross-modal Onset Difference Time (ODT).....	92
5.6 DISCUSSION.....	97
Chapter 6 Multimodal Onset Detection	99
6.1 ON MUSIC AND ONSETS	99
6.2 AUDIO VS. SENSOR ONSET DETECTION: STRENGTHS AND WEAKNESSES	102
6.3 SYSTEM DESIGN AND IMPLEMENTATION	104
6.3.1 Onset Detection Function.....	105
6.3.2 Fusion Algorithm.....	106
6.3.3 Data Collection	107
6.4 ONSET DETECTION AND FUSION RESULTS.....	107
6.4.1 Discussion: Audio-Only Onset Detection Results.....	108
6.4.2 Discussion: Sensor-Only Onset Detection Results	109
6.4.3 Discussion: Multimodal Onset Fusion Results	110
6.4.4 Discussion: Precision, Recall, and F_1 -Measure.....	111
6.5 MUSICAL CONTEXTS AND CONCLUSIONS	114
Chapter 7 Rethinking How We Learn: Performance Metrics and Multimodality in the Practice Room	117
7.1 BACKGROUND AND MOTIVATION	117
7.2 OVERVIEW OF METRICS EXPERIMENTS	120
7.3 DATA COLLECTION	120
7.3.1 Ezither Data	120
7.4 TEMPO METRICS AND STATISTICS	121
7.4.1 Tempo Estimation Algorithm	121
7.4.2 Tempo: Performance Timing	122
7.4.3 Tempo: Evolution of Timing over a Performance.....	123
7.5 BOW ARTICULATION TECHNIQUE METRICS AND STATISTICS	128
7.5.1 Definition of Bow Articulations.....	128
7.5.2 Bow Articulation: Tempo Accuracy	128
7.5.3 Bow Articulation: Onset Difference Time (ODT).....	135
7.5.4 Bow Articulation: Articulation Attack Slope.....	137
7.6 LONG-TERM METRICS ANALYSIS.....	139
7.6.1 Long-Term Tempo Metrics: Average.....	140

7.6.2 Long-Term Tempo Metrics: Standard Deviation.....	140
7.6.3 Long-Term Tempo Metrics: Range	144
7.6.4 Long-Term Bow Articulation Metrics: Tempo Accuracy	146
7.6.5 Long-Term Bow Articulation Metrics: Onset Difference Time	148
7.6.6 Long-Term Bow Articulation Metrics: Articulation Attack Slope.....	150
7.7 SUMMARY	152
Chapter 8 Conclusion	155
8.1 SUMMARY	155
8.2 PRIMARY CONTRIBUTIONS	156
8.2.1 Enabling Multimodal Musical Analysis with Nuance	157
8.2.2 Teaching the Computer to Know Who You Are.....	157
8.2.3 Negotiating Novel Understandings and Interactions in Drum Performance	158
8.2.4 Advancing Machine Musicianship through Multimodal Fusion	158
8.2.5 Refining The Way Musicians Learn: Multimodal Performance Metrics and Musical Pedagogy	159
8.3 PRINCIPLES AND CONSIDERATIONS ON THE DESIGN OF MULTIMODAL MUSICAL INSTRUMENTS AND SENSOR SYSTEMS	160
8.3.1 What is the musical context?.....	160
8.3.2 Exploitation	161
8.3.3 Transparency	162
8.3.4 Applying multimodality <i>to</i> a musical task vs. applying multimodality <i>into</i> a musical task.....	162
8.3.5 On Continuous Controls and ‘Leaky Faucets’	163
8.4 MAPPING MULTIMODAL MUSICAL SYSTEMS	164
8.4.1 One-to-one, Complimentary modalities, and multi-dimensional control	165
8.4.2 Many-to-one as a space for multimodal integration.....	166
8.4.3 Defining a set of parameterizations	167
8.5 FUTURE WORK	167
8.6 CONCLUSION	168
Appendix	169
Appendix A Live Performances and Applications	171
A.1 MINIM PERFORMANCE AT THE NEW ZEALAND SCHOOL OF MUSIC SONIC ARTS EXHIBITION CONCERT, OCTOBER 9TH, 2010	171

A.1.1	Excitation, Impulse, and Probability Machines	172
A.1.2	Composing by Improvisation	173
A.1.3	Performer Interaction	174
A.2	III: PERFORMANCE AT THE NEW ZEALAND SCHOOL OF MUSIC SONIC ARTS EXHIBITION CONCERT, OCTOBER 9 TH , 2011	174
A.2.1	Hyperinstruments and Gesture	175
A.2.2	Composition, Improvisation, and Iteration.....	177
A.2.3	Performer-Interaction.....	179
A.3	TRANSFORMATIONS: INTEGRATING MULTIMODAL MUSIC, DANCE, VISUALS, AND WEARABLE TECHNOLOGY, MAY 11 – 17, 2012.....	181
A.3.1	Designing the System.....	181
A.3.2	Discussion.....	185
A.4	SMARTFIDUCIAL.....	186
A.4.1	Background and Motivation.....	187
A.4.2	Implementation.....	188
A.4.3	Hardware.....	188
A.4.4	Software.....	192
8.6.2	Discussion: Spatial Relationships and Tangible Interfaces	193
A.4.5	Discussion: New Affordances for Tabletop Interaction.....	194
A.4.6	Final Thoughts on Augmented Fiducial Objects.....	195
A.5	SUMMARY.....	196
Appendix B	Sensors	199
B.1	TRANSDUCERS	199
B.2	PIEZOELECTRIC SENSORS	200
B.3	FORCE-SENSING RESISTORS.....	201
B.4	ACCELEROMETERS.....	202
Appendix C	Communication Systems and Protocols.....	205
C.1	MIDI.....	205
C.2	OPEN SOUND CONTROL.....	206
C.3	TUIO	207
Appendix D	Machine Learning	209
D.1	SUPERVISED LEARNING	209
D.2	K-FOLD CROSS-VALIDATION	211
D.3	ALGORITHMS	212

D.3.1 Decision Trees	212
D.3.2 Naive Bayes	213
D.3.3 k-Nearest Neighbor (kNN).....	214
D.3.4 Artificial Neural Networks (ANNs)	215
Appendix E Refereed Academic Journals and Publications	217
Bibliography	219

List of Figures

Figure 1: Example of the McGurk effect integrating /ga/ (visual) and /ba/ (auditory), results in the perceived /da/	3
Figure 2: Unimodal vs. Multimodal Musical Interfaces.....	6
Figure 3: Overview diagram of Complementary Modalities vs. Multimodal Fusion.....	7
Figure 4: Overview of Research.....	9
Figure 5: EMG biometric and gyro-based position controller (arm bands, headbands and base) used in (Tanaka and Knapp 2002)	17
Figure 6: Max Mathews and the Radio Baton (left) and the Buchla Lightning III (right)	21
Figure 7: Collaborative music making on the Reactable (left), and Bricktable “Roots” (right)	22
Figure 8: Final Hyperbow violin design by Diana Young.....	24
Figure 9: Esitar sensor systems, close up of thumb sensor (left), and usb, standard audio jack, knobs, buttons, and switches (right).....	36
Figure 10: Pictures of the Ezither hyperinstrument and bow	37
Figure 11: Picture of the Esuling controller showing the two FSRs and buttons.....	38
Figure 12: XXL sensor system (left) and screenshot of XXLSerial MIDI/OSC translator (right).....	39
Figure 13: Requirement comparison of other software and frameworks considered as of May 2012.....	43
Figure 14: Overview of Nuance input synchronization and output scheme	46
Figure 15: Audio Recorder object.....	47
Figure 16: Serial Message Format	47
Figure 17: Example Arduino serial out messages for two analog sensors.....	47
Figure 18: Example .xml configuration.....	48
Figure 19: Sensor (serial), OSC and MIDI Recorders	49
Figure 20: Nuance session main editor panel screenshot	50

Figure 21: Overview of the performer recognition system (only sitar shown in figure)	55
Figure 22: Overview of data capturing and feature extraction.....	58
Figure 23: Audio vs. sensor vs. multimodal accuracy achieved for improv data set after training with Exercise and <i>Yaman</i> data sets	64
Figure 24: Nuance software (Left) and custom sensor system (Right).....	65
Figure 25: Overview of drum rudiments and paradiddles performed by all performers for drum performer recognition in 4.4 and drum stroke computing in Chapter 5.....	66
Figure 26: Overview of features extracted at each event in the data set	67
Figure 27: Feature scatter-plots of audio features <i>regularity</i> vs. <i>roughness</i> on the left and <i>regularity</i> vs. <i>spectral centroid</i> on the right	69
Figure 28: Feature scatter-plot of audio feature <i>spectral rolloff</i> vs. sensor feature <i>average (mean) release phase deceleration</i>	70
Figure 29: Confusion matrix for all data sets and sensor features only using the MLP classifier	73
Figure 30: Performer recognition accuracy for all classifiers using all features and all data sets D1-D4	74
Figure 31: Accuracy for audio-only features vs. sensor features vs. multimodal features by averaging all classifiers	75
Figure 32: Overview of drum hand recognition system.....	82
Figure 33: Overview of onset detection algorithm	84
Figure 34: Average drum-hand recognition accuracy (%) across all performers for each classifier	88
Figure 35: Average classification of all classifiers for each performer	88
Figure 36: Classification results for the two best performing players when trained on all <i>other</i> data sets	91
Figure 37: Onset difference times for the 60-sec. of D1 (performer one top, performer two bottom)	92
Figure 38: Bar graph visualizing average onset difference time metrics (Table 13 - rush, lag, mean, standard deviation, and range) for all ten performers, in seconds.....	93

Figure 39: Snare drum waveform (left) and envelope representation (right) of the note <i>onset</i> (circle), <i>attack</i> (bold) and <i>transient</i> (dashed). Figure adapted from (Bello et al. 2005)	100
Figure 40: Strengths and Weaknesses of Audio and Sensor Onset Detection ...	104
Figure 41: General Overview of Multimodal Onset Fusion.....	105
Figure 42: Onset Curve (envelope) and peak-picked onsets (circles) for a short window of audio	105
Figure 43: Onset fusion algorithm pseudo-code	106
Figure 44: Audio onsets detected over an excerpt of mostly tremolo playing, TP (black vertical lines), FN (grey rectangles).....	108
Figure 45: Sensor (accelerometer) onsets detected over an excerpt of mostly tremolo playing, TP (black vertical lines), FP (grey rectangles).....	110
Figure 46: Multimodal fusion onsets detected over an excerpt of mostly tremolo playing, TP (black vertical lines), FN (grey rectangles).....	111
Figure 47: Comparison of precision, recall, and F_1 -Measure for audio, sensor, and fusion onsets	113
Figure 48: Comparison of TP, FP, FN and #Bows for Audio, Sensor, and Fusion Onsets	115
Figure 49: Melody Repeated in Data Set 3 (D3).....	121
Figure 50: Tempo estimated for three recordings from Ezither recording #7, data set D2.....	122
Figure 51: Tempo evolution of Ezither recording #7, data set D2, (a) andante, (b) moderato, (c) allegro	124
Figure 52: Box and whisker plot for Ezither recording #4, data set D1 showing bowing statistics of three bow strokes (detaché, martelé, and spiccato) when playing at the target tempo of 120bpm.....	129
Figure 53: Tempo estimate and statistics (mean, standard deviation, and range) for detaché, martelé, and spicatto bow strokes for the Ezither recording #4 data set D1 (target tempo = 120 bpm).....	131
Figure 54: Onset difference time (ODT) statistics for recording #9 data set D1	137

Figure 55: Note attack slope for Ezither recording #9 data set D1, Detaché entire recording (top), 2.8 second window from 10sec – 12.8sec (bottom)	138
Figure 56: Mean and standard deviation of bow stroke attack slopes for Ezither recording #9 data set D1 detaché, martelé, and spiccato	139
Figure 57: Average (mean) tempo of each D2 tempo recording (andante – bottom, moderato (middle), allegro (top), from the entire data corpus 1-16	140
Figure 58: <i>Standard deviation</i> for every session data set D2 tempo (solid) and linear trend lines (dashed)	142
Figure 59: <i>Range</i> for every session data set D2 tempo (solid) and linear trend lines (dashed).....	145
Figure 60: Standard deviation (A – top) and range (B – bottom) of bow articulation tempo across all D1 data sets collected.....	147
Figure 61: Session-to-session change in Ezither articulation Onset Difference Time.....	149
Figure 62: Ezither average articulation attack slope difference over time for (AAS difference – solid, trend lines dashed)	150
Figure 63: Ezither articulation attack slope standard deviation (top) and range (bottom) practice session-to-session difference over time for (AAS difference – solid, trend lines dashed).....	152
Figure 64: Curtis Bahn and the EDilruba (Sensor Esraj).....	176
Figure 65: Performance of <i>III</i> , group (left), and close-up of modified 12- string acoustic guitar and bow (right).....	179
Figure 66: Overview of dance technology, XXL accelerometers on hands, and Microsoft Kinect for real-time projection mapping (masking) onto the dancer	185
Figure 67: SmartFiducial System Overview Diagram.....	188
Figure 68: SmartFiducial hardware design and layout	189
Figure 69: SmartFiducial TUIO Protocol Specification.....	191
Figure 70: Overview of the SmartFiducial Serial Protocol	191
Figure 71: SmartFiducial Prototype (buttons 1 & 2 not pictured).....	192
Figure 72: Two SmartFiducials being used with Turbine	193

Figure 73: Overview of common forms of energy that transducers convert	200
Figure 74: Common FSR shapes and configurations (force a and b, position c)	201
Figure 75: Illustration of a typical supervised learning flow, adapted from (Fiebrink 2011).....	210
Figure 76: Sample decision tree constructed from the feature set of instruments (for genre classification) listed in Table 25	213
Figure 77: Illustration of kNN classifier where there are two classes (class 1 and class 2), and a rounded rectangle marked '?' in the center is the test or prediction point	214
Figure 78: Simple artificial neural network with one hidden layer.....	216

List of Tables

Table 1: Bol Patterns and Alankar exercises (data set 1)	57
Table 2: Accuracy achieved using audio only (15-second window)	60
Table 3: Accuracy achieved using sensors only (15-second window)	61
Table 4: Accuracy achieved using individual sensor features on all data sets, T=Thumb F=Fret (15-second window).....	62
Table 5: Accuracy achieved using multimodal data (15-second window)	62
Table 6: Identification accuracy of sensors vs. audio vs. multimodal fusion using a combined corpus from all data sets (at various window periods)	63
Table 7: Performer recognition accuracy using audio features (and ODT feature) only.....	72
Table 8: Performer recognition accuracy using sensor features only	72
Table 9: Performer recognition accuracy using all features (audio & sensors) combined	74
Table 10: Accelerometer onset detection accuracy for performers 1 and 2.....	85
Table 11: L/R Drum hand recognition accuracy for all performers and data.....	87
Table 12: Classification accuracy using separate rudiments for training and testing	90
Table 13: Average onset difference statistics for both performers.....	93
Table 14: Rush/Lag distribution of performer ODTs	96
Table 15: Distribution of onsets detected from audio-only as either True Positive, False Positive, or False Negative	109
Table 16: Distribution of onsets detected from sensor-only (accelerometer) as either True Positive, False Positive, or False Negative	110
Table 17: Distribution of onsets detected from multimodal onset fusion as True Positive, False Positive, or False Negative.....	111
Table 18: Comparison of precision, recall, and F_1 -Measure for audio, sensor, and fusion onsets.....	113

Table 19: Comparison of TP, FP, FN and #Bows for Audio, Sensor, and Fusion Onsets	115
Table 20: Tempo evolution statistics (min, max, mean, standard deviation, and range) of Ezither recording #7, data set D2, andante (80 bpm), moderato (110 bpm), and allegro (140 bpm)	123
Table 21: Five number summary for each bow stroke in Ezither recording #4, data set D1 (target tempo = 120 bpm).....	131
Table 22: One-way ANOVA multiple comparisons (Tukey HSD) for each bow stroke, Ezither recording #4 data set D1 (dependent variable = tempo).....	134
Table 23: Average Range of tempos from D2 for all data collected	145
Table 24: Tempo, range, and standard deviation averages over all practice sessions	146
Table 25: Sample feature set of instruments for genre classification problem using a decision tree classifier.....	212

Acknowledgements

This dissertation would not have been possible without the encouragement, direction, and inspiration of many people. First and foremost, I'd like to thank my primary advisor, Dr. Ajay Kapur. Ajay not only introduced me to the world of academia, and encouraged me to pursue postgraduate education, but he showed me that it is possible to simultaneously live in both academics and the arts. Among the many things I've learned from Ajay, he showed me it was possible to combine my passion for music and technology. His encouragement in developing my personality as a musician, engineer, and computer scientist, has been elemental over the course of my research. It goes without saying that Ajay has been and continues to be an inspirational source in my life, and I can't thank him enough.

Secondly, I owe endless thanks to my close colleague and friend, Owen Vallis. Working with Owen has been nothing short of amazing. His endless supply of brilliant ideas, creativity, and propensity to willingly dive head first into the deep-end and find his way out has always been a driving force, pushing me to achieve more. From peer coding to playing live music together as FlipMu, we're always on the same page, and Owen has been a hugely influential person in my life these past few years.

I'd also like to give many thanks to my second advisor Dr. Dugal McKinnon, and the New Zealand School of Music. Your encouragement has enabled me to achieve the work in this dissertation, your support has enabled me to focus on my research, and pursue the corners of my mind. Dugal, I am lucky to have received your perspective; from meetings over coffee, to your thought provoking posts on the Sonic Arts Facebook and email groups, I owe you many thanks. Additionally, I'd like to thank Victoria University and the Victoria Postgraduate Students Association, for supporting me with the Victoria PhD Scholarship, and the Postgraduate Research Excellence Award.

I also owe numerous thanks to the many individuals I've been lucky enough to have collaborated with over the last few years. To Dr. Matthew Wright, you have been highly influential, from your work on Open Sound Control, which has changed the way I interact with the computer musically, to collaborating with

you on sitar performer recognition in this dissertation. Your notes and guidance in what was one of my first research endeavors were crucial to my research approach, and was a springboard to the remainder of this dissertation. To Blake Johnston and Jason Erskine, it's been a pleasure working with you two and delving further into the world of hyperinstruments. You both brought your own voices to our practice sessions, our tech meetings, and uncovered incredible facets of composing and performing music with multimodal hyperinstruments. Blake, special thanks for bravely venturing into the world of Nuance, spending so much time working with the software, and letting me know when it worked, and when it didn't!

Lastly, it goes without saying that I owe the world and more to my family. Dad, Mom, Rami, Leyat, Natalie, and Nina—thank you for always supporting me, not only over the course of my PhD, but from day one. You have all inspired me in so many ways, and have made me who I am.