CARNEGIE MELLON UNIVERSITY

School of Architecture

College of Fine Arts

Thesis

Submitted in Partial Fulfillment of the requirements for the degree of

Master of Science in Computational Design

TITLE:

Drifting Through Design Research: An Autobiographical Rethinking of Constructing a Thesis

AUTHOR:

Yumeng Zhuang

ACCEPTED BY ADVISORY COMMITTEE:

Daragh Byrne	Principal Advisor	May 9th 2023 DATE
Lining Yao	Advisor	May 9, 2023 DATE
Sinan Goral	Advisor	May 9, 2023 DATE

Yumeng Zhuang

Drifting Through Design Research:

An Autobiographical Rethinking of Constructing a Thesis

SPRING 2023

Carnegie Mellon University
Master of Science in Computational Design

Committee: Daragh Byrne, Lining Yao, and Sinan Goral

Contents

How to Read This Thesis 15

Abstract 17

Forecast - Intimate Ambiguity: 11/23/2022 - 11/29/2022 18

Forecast - Hidden Labor in Machine Learning: 11/30/2022 - 1/27/2023 18

Forecast - Fallible Digital Memory: 2/7/2023 - 2/18/2023 19

Forecast - Tangible Rumination: 2/19/2023 - 3/17/2023 19

Journey Map 21

People 23

Introduction 25

Background 29

What is Design Research 29

How is Knowledge Generated 30

Drifting in the Process of Knowledge Generation 31

How are Design Research Documents Produced 33

Autobiographical Design 36

Ways to Narrate a Design Research Journey 37

Journey 41

The Start 41

Sensing Implementation 46

Generating Ambiguous Language 49

Landing: Intimate Ambiguity 52

Mapping Body Gestures to Intimacy 53

Discussion: Intimate Ambiguity 55

Control Implementation 57

Music Implementation 60

Landing: Hidden Labor in ML 65

Discussion: Hidden Labor in ML 72

Sound Implementation 73

Machine Learning Implementation 75

Landing: Fallible Digital Memory 77

Discussion: Fallible Digital Memory 85

Landing: Tangible Rumination 88

Pouch Implementation 91

Discussion: Tangible Rumination 97

Conclusion of the Journey 99

Discussion 103

Overview of the Discussion 103

Drifting from Physics 104

Drifting Categories by Krogh and Koskinen (2020) 105

Designing this Thesis 106

Types of Writing 107

Weaving 108

Formatting and Interstitials 111

Evaluating this Thesis 112

Drifts and the Practice of Documenting Drifts 113

When and How to Document Drifts 115

The Effectiveness of My Approach 116

Conclusions 119

Limitations and Future Work 121

Acknowledgements 123

Appendix A: Technical Implementation 133

Arduino Capacitive Sensing and Communication 133

Capacitive Sensing Characterizations 133

Slider Pouch 135

Slider Pouch Web Interface 136

Multi-Touch Pouch 137

Arduino Serial Print 139

Serial Communication to Python 140

Machine Learning 142

Regressor 142

Realistic Pouch Simulation 143

RL: DQN 145

Pitch Detection 146

ML Classification 150

Further Sound and Music Manipulation 152

Music Generation 152

Sound Segmentation 154

Sound Mixing 157

Wearable Prototype 158

Raincoat 158

Hat Making 158

Mask Making 161

Box Making 163

Appendix B: A Day of a Transparent Raincoat by the Window 165

Appendix C: Interview Questions 169

Discussion before Sharing Document 169

Shared Document and Questions 169

List of Figures

1 The chronological evolution of Journey concepts and implementations. 22

2 Pages from Equator Workbook 3, the third iteration of a design proposal, credit W. Gaver (2011). 34

3 Pages from the Critical Journal by Sadokierski (2019). 34

4 Annotations of the Photostroller, credit B. Gaver and Bowers (2012). 35

5 Part of page 18 of W. G. Sebald's *Austerlitz*, credit Sebald et al. (2011). 38

6 The original sticky pouch, developed with Violet Han and Di Wu in Professor Lining Yao's class in Spring 2022. 41

7 Left: Inflatable Mouse, credit Kim et al. (2008). Right: Neck Hammock, credit: Kidder-Wolff (2022) 42

8 The notes I took from reading the preface of *Designing with the Body: Somaesthetic Interaction Design* by Höök (2018) 43

9 Left: *exhale* by Schiphorst (2006). Right: The wearable from Squeaky/Pain, credit Demir et al. (2022). 44

10 Faraday Chair by Dunne & Raby (1995). Photo credit: Lubna Hammoud 45

11 Left: the slider pouch in action. Right: the back of the pouch, with a uniform rectangular electrode. 47

12 The simpler pouch with crossing electrodes. 47

13 The Arduino analogRead response to a pinch on the electrodes. 48

14 Edward Lear's limerick: There was a Young Lady of Portugal. 49

15 Three lists of words I thought about using to generate ambiguous language. 50

16 The components of the proposed system: mechanized on–body pouches with sensors feeding into a computer that generates related ambiguous language to the wearer. 51

17 A brief digression to collaborative poetry. 53

18 Attempts to draw a few gestures. 54

19 Some thought and emotions are mapped to different body parts and some thoughts on output modalities. 54

20 Fantasy: "I love your shoulder." 55

21 Collection of ML schemes, with different input and output modes. 57

22 Output of a simulated pump with ML regression of pump speed. The orange and blue curves are pouch volume curves as the results of the pump speed. 58

23 Actions and effects the human operator and the machine pump could take. 59

24 Hand-calculated PV curves and the transitions between different pouch states. 59

25 Calibre, credit Jang (2013) 61

26 Open Surgery, credit Kolkman (2015) 62

27 (author)rise, credit Agrawal et al. (2018) 62

28 Pero sans Cimon, credit Bourgault et al. (2020) 63

29 Created by me and DALL-E (https://openai.com/product/dall-e-2). 64

30 Brainstorming for concepts related to ML and labor. 66

31 Left: visualization of the training process. The music notes are displayed above, and orange denotes the note I was singing, and green denotes the note the machine was predicting. Values are note numbers divided by 6. Right: The test setup of the pouch. 67

32 What I showed to Professor Jesse Stiles. 70

33 Audio segmentation. Box 1: identified sequence 1; Box 2: identified sequence 2; Box 3: sound interference. 73

34 Concept mapping stemming from the poppable plastic bag. 74

35 Illustration of how audio data and capacitive sensing feed into the ML model. 75

36 The Memory Device, credit Bertran (2014) 77

37 Components of the Transparent Puffer Coat. 78

38 A few wearable precedents that inspired the development of my wearable. a) Sculpted Cashmere, Fang (2013), photo credit: Quinn Connors b) Air Slaves, Let it Be! art agency (2017) c) mediated reality, ling QL (2013) d) Intangible Embrace, Eriksson (2011) e) Freakshow, Birch (2013) f) FINAL HOME, Tsumura (2015) 79

39 Various forms that were inspired by the precedents. 80

40 The experiment to sew pouches onto a crochet piece. 81

41 The experiment with a pouch under my sweater. 81

42 How the sensor strips are laid out on the plastic raincoat. 82
43 Singing the phrase with the arm raised. 82
44 Illustration for the fictional "use-scenario". 83
45 The comic that accompanies the raincoat story. 84
46 The concepts around body and sound which led to reflection and rumination. 86
47 An enriched and refined version of concepts leading to the "here-and-now". 86
48 The curling tail pouch. Left: the pouch was pinched. Middle: the pouch was relaxed. Right: the seal pattern of the curling tail. 92
49 Taping one curling tail to my forehead and overlaying and mirroring many photos. 93
50 Some early thoughts on head-mount wearables. 93
51 The (disappointing) process of making an inflatable hat. 94
52 What the "hat" looked like on someone's head. 94
53 The Echo Mask in its current form. 96
54 Some core elements I wanted to keep that I developed while working on Hidden Labor in ML. 98

55 A shorter (and therefore bigger) interstitial. 112
56 A longer interstitial. 112

57 The original sticky pouch, developed with Violet Han and Di Wu in Professor Lining Yao's class in Spring 2022. 133
58 The circuit and the post-it notes used to characterize MTK signal. 134
59 First row: button pouch: single point contact detection via capacitive sensing. Second row: slider pouch: linear contact detector. Last two rows: a push-pull button with 4 states. 135
60 First row: button pouch: single point contact detection via capacitive sensing. Second row: slider pouch: linear contact detector. Last two rows: a push-pull button with 4 states. 136
61 The characterization of the slider pouch fits the theory of capacitive sensing. 136
62 The web screenshot of the slider pouch with its corresponding state. 138
63 How a person interacts with a single pouch and it sends data to Python. 138

64 Left: the Arduino setup for sensing the pouch inputs. The black board is Elegoo Uno, the blue board is a multiplexer (CD74HC4607). Right: The bottom four crossing lines are the multi-touch electrodes used in this experiment. The top pattern is another option. 138

65 The circuit diagram from the Multi-Touch Kit Tutorial. Credit: Pourjafarian et al. (2019). 139

66 Output of a simulated pump with ML regression of pump speed, with l2 regularization and small gaps between acceptable ranges. The orange and blue curves are pouch volume curves as the results of the pump speed. 142

67 Earlier attempts to model the pouch dynamics with straight lines based on its states. 144

68 A more sophisticated model of the air in pouches, taking into account air inflow. 144

69 Output of a simulated pump with Reinforcement Learning of pump speed. The blue curve is the pouch volume curve as a result of the pump speed. The orange curve is the reward the agent receives for being within the pink target ranges. 146

70 A web-based training visualizer with annotations. 151

71 The annotated Max Patch that mixes the recording clips based on incoming relative weights. 159

72 The raincoat with crossing copper strips. 160

73 Left: the paper pieces were taped onto a hat. Middle and right: the pieces prepared to be made into pouches. 160

74 Left: the fitting goes through one layer of plastic. Right: the setup for sealing the pouches. 160

75 Left: how I used the T connector. Middle: the hat on the model. Right: the hat on its own. 161

76 The shape and the bag I used to make the Echo Mask. 162

77 Left: the anatomy of the four pouches and the layout of their sensors. Right: how the sensors make contact under a pinch. 162

78 Some of the cut lines for the box. 163

79 Annotated components of the Echo Mask. 164

Glossary

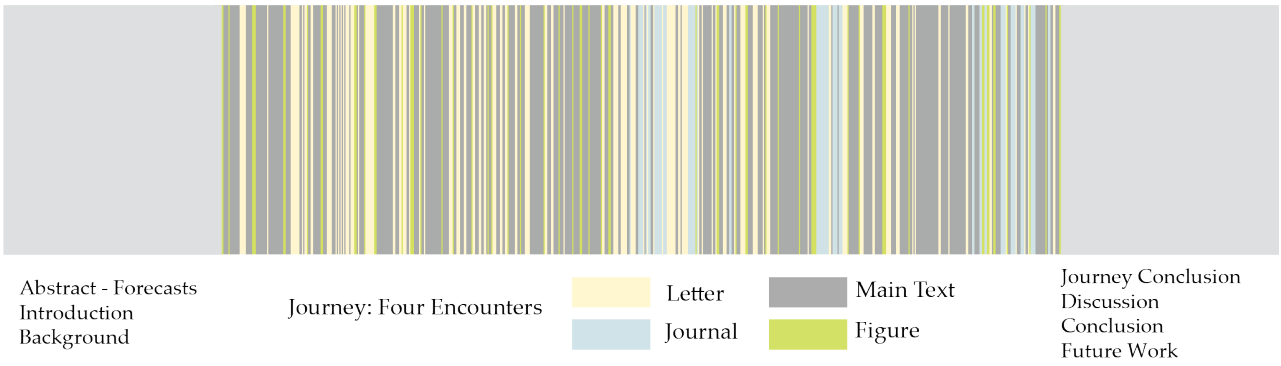
autobiographical of, relating to, or influenced by one's life or past personal experiences (Merriam-Webster Dictionary). 17

design research An activity of knowledge generation through acts of design or making, which may involve constructing objects, concepts, or theories. 17

drifting An unplanned process or event that brings a research project away from its original concept or approach, analyzed in detail by Krogh and Koskinen (2020). 17

Plate tectonics explains why Earth's continents are moving; the theory of continental drift did not provide an explanation. [...] This shift from one theory to another is an example of the scientific process: As more observations are made and measurements are collected, scientists revise their theories to be more accurate and consistent with the natural world. National Geographic Society (2022).

How to Read This Thesis



The following materials are roughly divided into three parts. First, this document begins with a standard academic form including Abstract, Introduction, and Background chapters. The body of the thesis is written in the form of a continuous Journey, in which I interweave three styles of writing along with figures. The *letter* style was written in a continuous stream of narration, shifted to the right in a lighter shade of gray. The *journal* entries are excerpts from my research journal in a typewriter font in a centered narrow column. The *main text* is using the tone of standard academic writing, occupying the left column. There are four major encounters scattered along the Journey, each of which has a landing section dedicated to it in the format of an introduction to a smaller imaginary thesis for that encounter. There is also an end part of the thesis. Conclusion of the Journey wraps up my explorations in this thesis, and presents my findings. Discussion, Conclusions, and Limitations and Future Work address this thesis as a whole.

Abstract

From the fall of 2022 to the present moment (the spring of 2023), I have been conducting thesis research on an interactive inflated air pouch with a sticky interior, looking for its purpose and a potential line of discourse to formulate around it. This design research project has opened up my mind to alternative ways to produce knowledge beyond the scientific and engineering research methods, which I was trained to apply prior to enrolling in this degree program.

During the development process, I iterated on the interactive structure, the air pouches, as well as numerous versions of secondary artifacts that supported the interactions. Told from my autobiographical perspective of making this thesis, I the concept of drifting in design research and walk the readers through my journey of developing and framing the interactive inflatable structure. The journey is scattered with mind maps and non-linear paths toward key inspirations, related literature background, documentation of the construction, casual and formal reflections, interactions with people who supported me, and self-evaluations. Interweaving different styles to tell the story, I provide a holistic view of my experience learning to conduct design research and focus on revealing the drifts in the journey. Through a detailed evaluation and conversations with students adopting design research for their theses, I propose that this alternative way to document design research could reveal the often hidden drifts in the process, bring readers closer to how knowledge is produced in full reality, and challenge existing norms of disseminating knowledge.

There are four major drifts in the propositions along the journey. I call the four resulting framings of my research *encounters* and provide the associated imaginary abstracts as Forecasts of them in the upcoming Journey.

Forecast - Intimate Ambiguity: 11/23/2022 - 11/29/2022

Bodies are our most intimate environment, yet our feelings and attitudes towards our bodies fluctuate. We express such fluctuations with ambiguous gestures, in front of mirrors, in the showers, but also with internal dialogues with ourselves. This thesis proposes an interactive wearable that maps ambiguous body gestures to ambiguous language about intimacy and examines what this mapping could offer to one's growing attitude towards self. The wearable is made out of deformable air pouches with sensors that read the shape information of the pouches. The signals from the sensors are unstable, which symbolizes the ambiguous gesture of the wearer towards themselves. Through a natural language generation system, the unstable signals are mapped to ambiguous poetry about bodies. The poems are then played back to the user, creating an illusion of an internal dialogue. This thesis provides autobiographical accounts of interacting with this wearable device at home and in a community space and discusses the implications of tangible devices that participate in intimate life.

Forecast - Hidden Labor in Machine Learning: 11/30/2022 - 1/27/2023

The propagation of Artificial Intelligence (AI) in many aspects of life is phenomenal. The abilities of intelligent machines improve, as they have taken advantage of more and more data that capture human life. Such data, especially at their early stages, are often crowdsourced and manually prepared in ways such as manual labeling. Users of intelligent machines are often unaware of the intensive labor involved in developing the products they are using and remain detached from the resulting intelligence. This thesis proposes an exploration of how to make the data collection and training process of machine learning (ML) tangible and creatively expressive, in order to bring the hidden labor in ML in contact with an individual's body and mind.

InVest is a wearable system that requires frequent inputs and corrections from the wearer to work as a companion singer for the wearer. The wearer would define a set of gestures by manipulating the sensor interface on the vest while giving a vocal cue, and the machine would learn to associate the cue with the sensor inputs and produce the sound when the same gesture reoccurs. It takes time for the machine to learn to classify different gestures apart, and the gestures also shift during the interaction. On the other hand, the wearer may forget the previous cue and has to train the machine

again. This collaborative effort presents a lively, noisy, and possibly frustrating performance.

This thesis examines the aesthetics of training an intimate machine through the user's body and reflects on the dialectics of using depersonalized effort versus using individual resources to generate technology. InVest could be a starting point for conversations around AI.

Forecast - Fallible Digital Memory: 2/7/2023 - 2/18/2023

Common digital devices such as phones and cameras are excellent at precisely capturing memorable events with almost no data loss at the time of retrieval. This thesis examines an alternative way to communicate and store information that disrupts the expectations of a memory storage device. The proposed device is forgetful and uncertain, yet malleable and unique to the user. I propose the design of the Transparent Puffer Coat, a wearable system that you can store in and retrieve from sound-based memory. It uses machine learning to register the shape deformation of the Coat to the sounds and reproduces the sound later when it returns to a previously registered deformation. I reflect on how living with the Coat changes my perception of what is worth recording, how I personify the Coat through my interaction, and how the distortion of memory playback reinforces my own memory. Through autobiographical studies, this thesis identifies design opportunities in the interplay between machines, bodies, and remembering.

Forecast - Tangible Rumination: 2/19/2023 - 3/17/2023

We expect our digital devices to be able to record and recall our spoken thoughts as in voice memos, for example. However, there are moments when we ruminate in our heads in purposeless, looping manners. This thesis presents the Echo Mask, an inflatable mask that challenges the expectation above and provides an alternative relationship between users and digital devices, in which rumination is mediated by reducing the meaning of thoughts.

The user interacts with the Echo Mask by first repeating a phrase they want to record and manipulating the mask into a specific shape. The algorithm pipeline with a machine learning model then registers the shape with the phrase, allowing the user to retrieve the phrase by reconstructing the shape of the mask they provided. The retrieved sound is played back with all the other recordings in a mix according

to the uncertainties in the output of the machine learning algorithm. Through the repetition and the chaotic playback, the user experiences a loss of meaning of the phrase they provide.

This thesis provides the implementation of the Echo Mask and the discussion about its implications for alternative digital devices. I argue that designing for rumination is a rich space that could encourage us to live in the moment and rethink what interacting with digital devices brings to our cognitive processes.

Journey Map

Figure 1 offers a flowchart that illustrates how the concepts and the implementation evolved and drifted during the development, as well as the forces that drove the drifts. As in the forecast, the four encounters and their corresponding Landing page in the Journey are linked here: *Intimate Ambiguity*, *Hidden Labor in ML*, *Fallible Digital Memory*, *Tangible Rumination*.

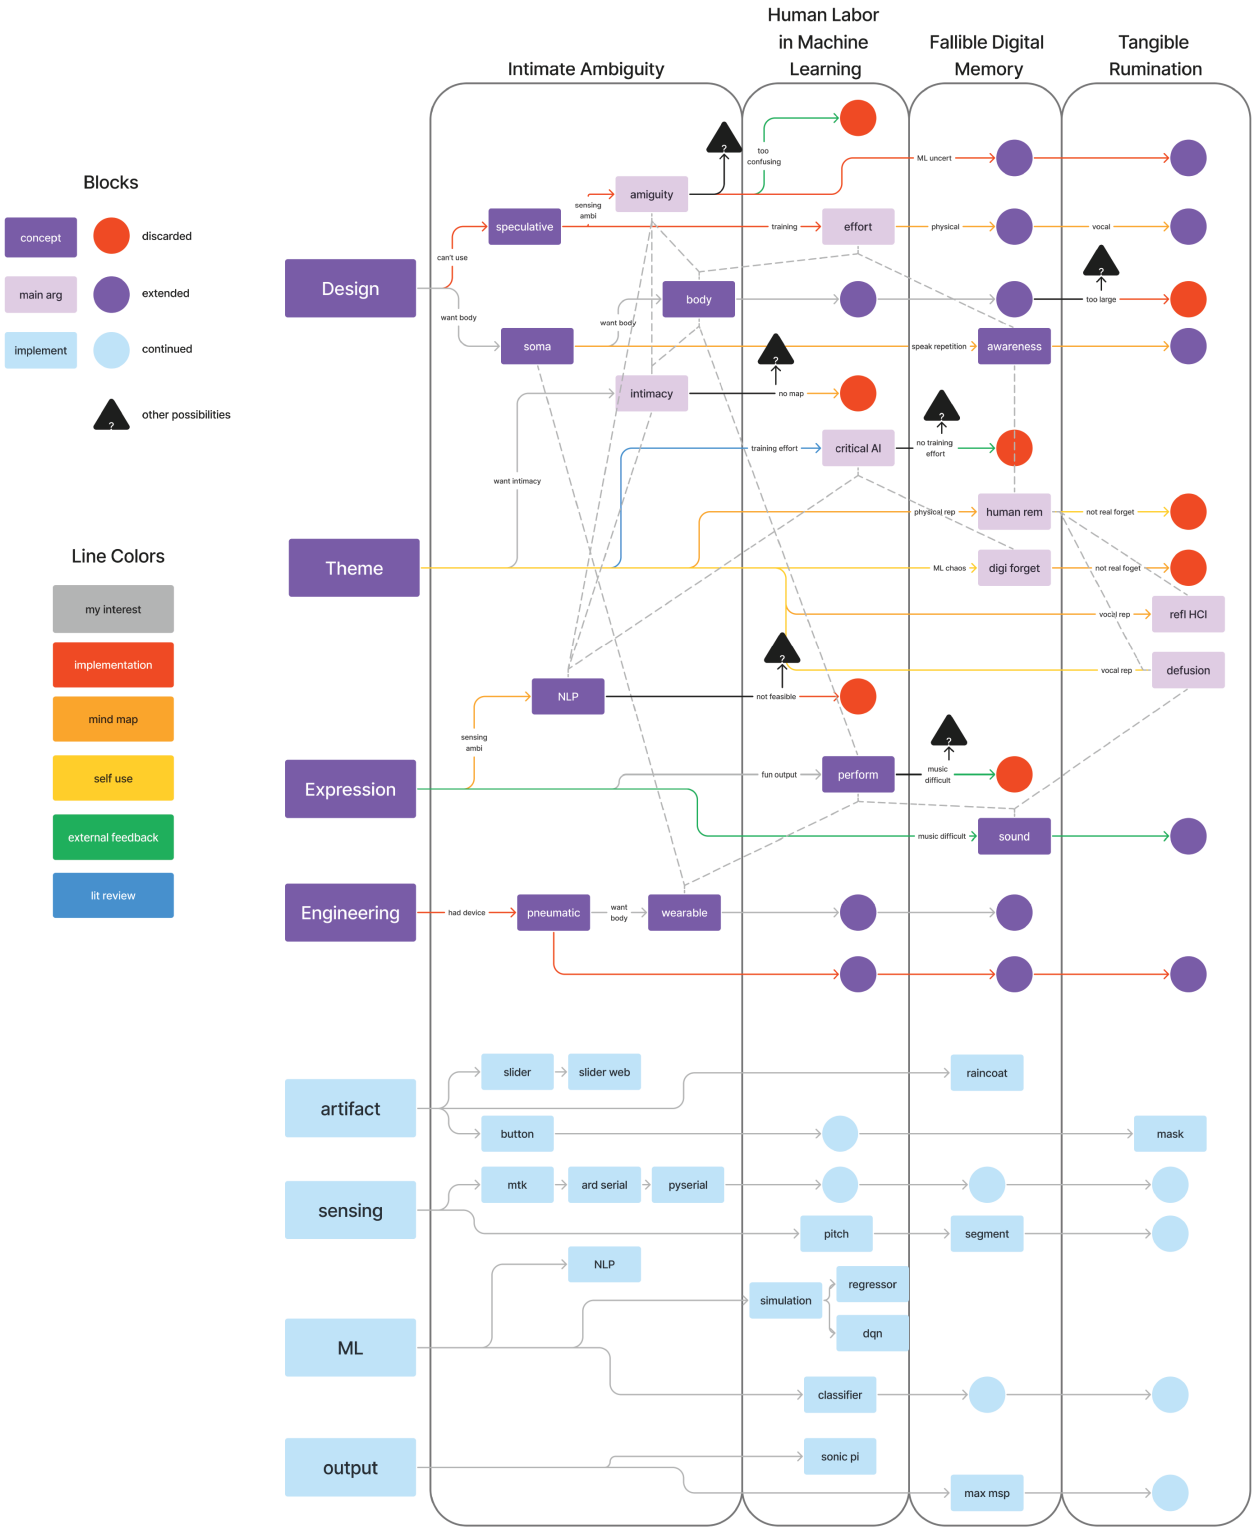


Figure 1: The chronological evolution of Journey concepts and implementations.

People

Here I provide the names of the people who have given feedback or otherwise influenced the trajectory of this thesis. The first group of people is my advisors, the second group contains other faculty members, and the last group lists many colleagues, friends, and a family member who have been part of the journey as well. I thank them for their contributions formally in the Acknowledgements.

Full Name	Subsequent Occurrences
Daragh Byrne	D. Byrne
Lining Yao	L. Yao
Sinan Goral	S. Goral
Marti Louw	M. Louw
Jesse Stiles	J. Stiles
Zhenfang Chen	Z. Chen
Yanwen Dong	Y. Dong
Howard Han	H. Han
Violet Han	V. Han
Malika Khurana	M. Khurana
Jingyang Liu	J. Liu
Hongyu Mao	H. Mao
Tomas Cabezon Pedroso	T. Cabezon Pedroso
Di Wu	D. Wu
Yanli Xie	Y. Xie
Chengzhi Zhang	C. Zhang

Introduction

Students from universities around the world are often required to complete some forms of documents that contain their research work when they graduate. In the numerous subfields of research, this thesis is concerned with design research, a form of knowledge production through acts of design and making, which may involve constructing objects, concepts, or theories. As a student transitioning from a purely scientific discipline to attempting to do design research work, I have encountered numerous challenges adapting to this form of research that ended up in "drifts", or, "those actions that take design away from its original brief or question and lead to a result that was not anticipated in the beginning (Krogh & Koskinen, 2020)". The amount and intensity of the drifts inspired me to analyze and discuss their implications for disseminating knowledge.

Specifically, my thesis raises the following questions:

- What are the common types of drifts in design research — specifically in design research theses — and what effects do they have?
- How can design research reconcile with drifts and reveal them constructively?

My thesis attempts to address these research questions through a documentation experiment of a design research project. This project involves making and framing the arguments around an inflatable pouch with a sticky interior and the secondary physical components and algorithms that support the interactions with the pouch. This underlying journey serves as a device for the analysis of drifts.

I noted that several factors prompted major drifts in the development of the artifact and the discourse: my personal experiences and reasoning, the external advice and feedback I received, and the material and technical affordances discovered through implementation.

The drifting resulted in roughly four major encounters in the iterative journey: First, in *Intimate Ambiguity*, I considered the sensing ambiguity of the pouches to be the dominant feature and argued that

it could be used to mediate intimacy. Second, in *Hidden Labor in ML*, I designed the machine learning (ML) component of the pouch system to be a potential tool to reveal human labor in training ML systems. Third, in *Fallible Digital Memory*, I rendered the recorded information of the pouch-sound system as a metaphor for storing fallible memory in digital devices. Fourth, in *Tangible Rumination*, I bridged the tactile feedback of the pouches with the sound quality of the output system and landed on how they may act as a rumination companion.

In the Journey chapter, I narrate the research process in three writing styles. The *letter* style addresses the reader and provides insights and reflections in a continuous stream of narration. The *journal* entries come from my research journal I kept to capture intimate moments of development. The *main text* is academic writing that provides grounding information about the encounters as well as technical information and literature background. I interweave these three styles densely to create a rhythm, and include original sketches and mind maps to enrich the reading experience.

Unlike many autobiographical design projects that produce a tool or a system, this thesis uses the methods of autobiographical design in an uncanonical way since I consider this thesis as a design artifact itself. As summarized by Desjardins and Wakkary (2016), "successful autobiographical design needs to fulfill a genuine need, the system needs to be lived with over a long period of time, and the designer needs to be the user of the system". **Hence, this thesis is about the process of constructing design research theses rather than a collection of unfinished theses I have made.** My genuine needs are beyond the need to have a thesis to graduate, but rather the passion to explore who I am and what I could offer to the design research community. I lived with the thesis, and the interactive system it centered on, by generating multiple forms of documentation and engaging in dialogues with colleagues, advisors, and external critics. The way I "use" my thesis is through constant development and evaluation, as well as writing.

The main contributions of this thesis lie in the following areas:

- This thesis offers a richly documented account and case study of drifts in design research.
- This thesis illustrates several sources and types of drifts that can be found in design research.
- It provides an alternative method to document the drifts that acknowledges and celebrates them.
- This documentation method also brings the readers to the core of how knowledge is generated in design research and invites them

to participate in the process by offering multiple narration styles.

- This thesis calls for innovations from the community in design research documentation.

In Evaluating this Thesis, this thesis is assessed based on self-reflection on the writing tone and quality and how they match the following criteria: critical, reflective, intimate, explorative, and receptive.

- **Critical** of what research approach and drifts were happening;
- **Reflect** on drifting moments and motivations;
- Provide **Intimate** accounts that reveal thoughts and feelings of the researcher to the reader;
- **Explore** modes of storytelling within a text-based thesis format; and,
- **Receptive** to other ways of doing a thesis.

I also evaluate this thesis by having conversations with students who adopt design research methods for their theses. I found that:

- Drifts are very common.
- Drifts are underdocumented in their writings, while writing about drifts is valuable for other researchers.
- Writing about drifts also places more focus on the research process instead of the concept or argument.
- My writing format and styles are effective in bringing the reader into what I thought and felt while doing research and the letter style provides a lighter and easier read.

The next chapter, Background, provides the background literature for design research, ways to document design research, and autobiographical design as a method. The following chapter of the thesis is the Journey through the four major encounters of framing the technological artifact: *Intimate Ambiguity*, *Hidden Labor in ML*, *Fallible Digital Memory*, and*Tangible Rumination*.

During the journey, each time I drift and land on a new framing for the project, I present the literature that informed the argument and how that converses with my existing knowledge and interests, the argument itself, the technical implementation to support the argument, and the evaluation of that implementation based on first-person experience and external feedback.

After the Journey and in the Discussion chapter, I summarize the potential drifting points and provide a critical reflection on how they

shaped the journey. I discuss how the three writing styles collaborate to convey drifting effectively. An evaluation of the thesis described above will be provided at this point. The Conclusions chapter summarizes the key findings of this thesis. Lastly, I open the ground to encourage future students and faculty to innovate in documentation methods, especially for interdisciplinary inquiries.

Background

What is Design Research

Design Research, a process that is the focus of this thesis, is an active field of research that has a myriad of definitions. I consider the following definitions as interchangeable terms.

Krogh and Koskinen (2020) use Koskinen's coined term *Constructive Design Research* to distinguish design research that generates knowledge "through constructing" from those design practices that stop at constructing artifacts.

Similarly, for Zimmerman et al. (2007), *Research-through-Design* is defined as "an inquiry focused on producing a contribution of knowledge" and is different from the practice in the Human Computer Interaction (HCI) community where design research is often used to aid product development directly. Research through Design is often conducted by designers, defined as people with practical training in design disciplines, such as industrial design and architecture.

B. Gaver et al. (1999) distinguish *Design as Research* from various related disciplines. They have a focus on Desgin as Research on new technologies. From the angles of artists and designers, they consider Design as Research different from typical research in science and engineering because Design as Research does not focus on applying the correct methods or making accurate and scientific analyses. Design as Research is also different from traditional design that lands on commercial products or meets user needs. Instead, Design as Research as covered in their article generates cultural, sociotechnical, and aesthetic insights into design "outside the norm" as well as new experiences and spaces.

Frayling (1994) divided Research in Art and Design into three categories: research into art and design, research through art and design, and research for art and design. The first category describes research into the historical, theoretical, and aesthetic aspects of art and design,

such as art history, as a form of knowledge production. The second category describes research that achieves communicable results through making and doing. The third category presents research that ends in an artifact, which embodies the knowledge and thinking in its existence. I consider the last two categories fit for the definition of Design Research that "constructs" and produces knowledge.

How is Knowledge Generated

In *Drifting by Intention: Four Epistemic Traditions from within Constructive Design Research*, Krogh and Koskinen (2020) categorize the sources of knowledge into four categories: the correct application of rational methods, the artifact itself, the discourse around the artifact, and insight gained through user participation and co-design processes. I am borrowing this classification to survey how knowledge is generated in design research.

The first category is similar to the research practice in science and engineering disciplines. It is a valuable proposition to locate knowledge in rigorous methods, since the methods could be applied across designers and foster an agreement on whether a design research project is valid or not. For example, to better understand how current HCI practices support everyday reflection, Mols et al. (2016) designed a web survey with multiple choice questions, checklists, and open-ended questions, and analyzed the data with statistical methods such as ANOVA.

The second category locates knowledge in the artifact. This placement is similar to Frayling's research for art and design (Frayling, 1994). Carroll and Kellogg (1989) raised the concept of artifacts as "theory-nexuses" where the theoretical development of HCI artifacts is embodied in the artifact themselves. According to the authors, hermeneutics, or the interpretation of the research artifact, could lead the theoretical research, and technological advancement could flow from the artifact to the theory. Zimmerman et al. (2010) also observed in their interviews that theories, which I consider as a form of knowledge as well, could result from analyzing the artifact to "reveal underlying relationships". For example, the Desktop Odometer is a device that shows the cumulative distance between all servers the user's internet queries pass through, which is designed to bring awareness to the hidden physical infrastructure that enables today's internet (Viny & Desjardins, 2020). In this artifact, the message is embedded in its functions. When the authors released this device to the market, they received new feedback from customer reviews that touched on climate change, a topic it was not explicitly designed to

engage.

The third category grounds knowledge in the framework and discourse built around research through design projects. Zimmerman et al. (2010) described several kinds of theories that could result from research through design beyond the study of the artifact. Theory on design, coming from studying the design process, produces knowledge about how designers work and why in order to improve design as a practice. Guiding theories to design work could also be an outcome of design research, which designers use to frame and reframe the design problem. Design implications that open up conversations on future design work and changes to be made to existing designs can also be valuable knowledge. At a broader level, Krogh and Koskinen used the concept of "program" from architecture to describe that knowledge can be generated through a lineage of frameworks and theories from a series of studies with one informing the next (Krogh & Koskinen, 2020). For example, Green (2019) raises an argument that encourages interaction designers to look for inspiration in "strange" places. The author found inspiration in Giorgio de Chirico's painting and designed a pair of interactive furnishings that enable independent living. During the development, the author also drew more inspiration from two contemporary dance performances to refine the product. It is through the design process that the author raises the argument that designers of tangible, embedded, and embodied interaction could find inspiration in domains far from HCI.

The fourth category argues that knowledge and insights can be drawn from users' experience and feedback through processes such as co-design (Krogh & Koskinen, 2020). Researchers learn from people and direct their design with the knowledge "distilled" from user experience, which makes their knowledge disputable instead of the rational and logical knowledge set in stone in science and engineering that results from the rigorous application of methods.

Drifting in the Process of Knowledge Generation

Drifting by Intention: Four Epistemic Traditions from within Constructive Design Research by Krogh and Koskinen (2020) is the book that inspired me to study explicitly how the framing and artifact of my thesis research shifted throughout the development of the interactive structure.

In all four systems of knowledge generation discussed in the previous section, Krogh and Koskinen identified ways in which knowledge drives drifting in research. The four systems of knowledge generation are called epistemic traditions.

In the Methodical Tradition, which corresponds to the first category above, drifting happens when there is explicit and rational evidence or knowledge, often generated in the early stages of the development, that warrants a change in research direction or hypothesis (Krogh & Koskinen, 2020). No personal whims or tacit knowledge could justify these drifts. This kind of drift happens more often in the traditional HCI literature, however, I find it often untraceable in published works. One example maybe be from Pourjafarian et al. (2019), where they first looked for the sinusoidal signal that maximized the difference between the touched state of the sensors and when the sensor was released. They documented that for commercial microcontrollers, it was difficult to generate pure signals at the desired frequency, so they turned to PWM signals that had the prominent component of that high frequency. This drift was necessary because of the hardware limitations.

In the Experiential Tradition, the artifact's affordances as well as the design process could drive the change in the knowledge generated (Krogh & Koskinen, 2020). Design researchers get more freedom to modify their research based on their experience and intuition. When Liu et al. (2018) was developing the Spore Stepper, she only intended to design a walking stick that collected soil that contained mushroom spores. However, using the walking stick in the field revealed to her that she might be collecting other forms of life, and the feeling of a heavier stick with each sample collection was fascinating and possibly unexpected as well. She proposed that future iterations on this device would need to address the weight of the stick. These discoveries and future directions would not exist if there were no artifact or no experience using it.

In the Programmatic Tradition, multiple frameworks and theories produced by design researchers sometimes need to be evaluated and decided on which ones to pursue further or to modify (Krogh & Koskinen, 2020). When design researchers interpret their works and form discourses around them differently, this kind of drift could occur. Desjardins and Key (2020) provide an honest account of a major drift in their research on interacting with data from homes, when two new members of the team collected and presented existing examples of what they had wanted to achieve. Upon taking a new look at existing efforts, the team found it not fitting their "ethos" and turned to investigate the nature of data. Here, the group evaluation of their work prompted the drift.

In the Dialectical Tradition where knowledge is produced through learning from user experience, drifts could occur when the learned results misalign with the design researcher's assumptions (Krogh &

Koskinen, 2020). The need to push the research project forward often results in drifts that balance the researcher's goals and the users' needs. Again, Desjardins and Key (2020) include their experience of designing open-ended probes for participants to speculate on how the device might work as a sensor in their homes, but they only received one response that did not meet their expectations. This attempt proved to the researchers that the effort devoted to designing these probes was not very helpful.

Desjardins and Key (2020) document more accounts of how their team approached a Research through Design project through the metaphor of different line shapes and qualities. The first example above comes from "perpendiculars" and the second example comes from "tangents". The rest of their analysis covers many stages of Research through Design beyond the concept of drifting, such as moving through uncertain areas, hitting dead ends, and conducting concurrent work. This thesis, however, is mostly concerned with the concept of drifting.

How are Design Research Documents Produced

I identify three goals for which documents are produced in the process of design research: The first aims to *gather materials* for disseminating the produced knowledge in the form of publications; Second, some documents are for publication, made available as contributions to the research community; the third goal is often called documentation, which is more than gathering materials but also supports the development of the design research *at the moment*.

Drifting is evident at various levels in these three types of documents as indicated by their purposes. Ordered by the stage of research where these documents are generated, I discuss five types of documents that are different from traditional research papers here: design workbook (W. Gaver, 2011), a Word document example from Bardzell et al. (2016), Critical Journal (Sadokierski, 2019), Project Reflection Tool (Dalsgaard & Halskov, 2012), and annotated portfolio (B. Gaver & Bowers, 2012).

A design workbook from W. Gaver (2011) is not a journal-like workbook, but rather a design proposal where multiple open ideas are documented and presented. The document may include text and images from many sources that present ideas and approaches that are "still malleable to change", as in Fig. 2. The possibility to change can be considered to be the possibility to drift as well. The design workbooks, however, stop here. They only support the early stages of design research. Although the author notes that these workbooks

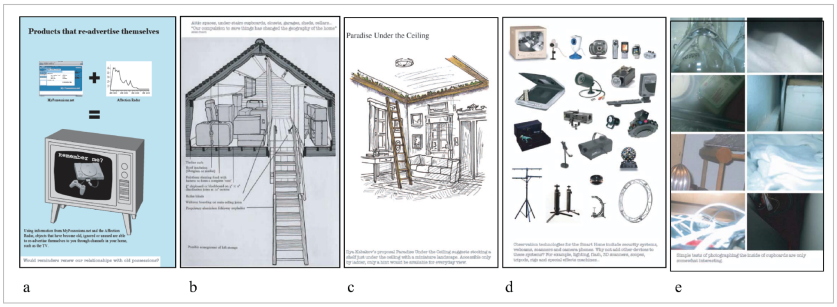


Figure 2: Pages from Equator Workbook 3, the third iteration of a design proposal, credit W. Gaver (2011).

contain "explorations of topics and approaches that extend beyond the eventual outcomes of design project", he does not elaborate on this point and implies that the eventual outcome does not include this workbook.

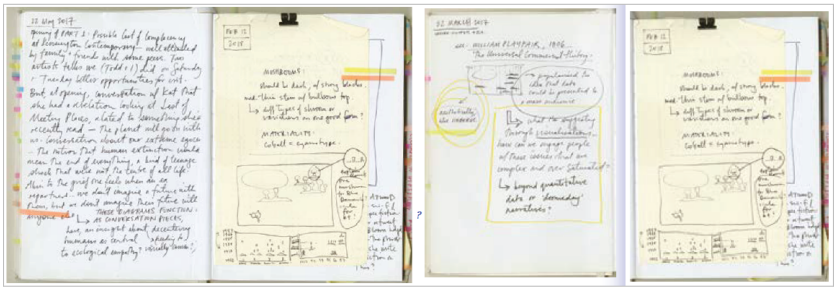


Figure 3: Pages from the Critical Journal by Sadokierski (2019).

The Word document example from Bardzell et al. (2016) and the Critical Journal approach by Sadokierski (2019) are similar documents developed throughout the research process that both support the development of design research as well as contain raw materials to draw from for publication. They typically contain descriptions of the process, research questions, images and even videos for digital documents, technical details, and so on. They evolve as the project progresses, therefore faithfully capture the drifts and decisions made. They also contain a significant amount of reflections on which researchers build further inquiry. The Critical Journal (for example, in Fig. 3) is more edited and annotated after the research project than the Word document example and is sometimes turned in as coursework as instructed by Sadokierski (2019). However, the materials from these two types of documents are only indirectly included in the final publication.

Another source of documentation that captures drifts is the Project Reflection Tool (PRT) (Dalsgaard & Halskov, 2012). This tool was designed to explicitly capture "design events" and prompt reflections

in the research process. It also records text and multimedia entries associated with events on a timeline. These materials are often helpful when the researchers move to the phase to write the publication. Although in-the-moment reflections were designed to inform the research process, Bardzell et al. (2016) found that researchers in their example case often logged the event after, not during, the event, and often left the reflection blank. Therefore, although drifts are evident in the events, they are not used to support the research as much as designed by the tool.



Figure 4: Annotations of the Photostroller, credit B. Gaver and Bowers (2012).

Lastly, the example I include here that appears after the research process that is somewhat unique to design research is the annotated portfolio (B. Gaver & Bowers, 2012). This is different from the traditional, text-heavy paper format in that it contains images with only succinct textual annotations (Fig. 4). These annotated images often present gathered relationships among a collection of design objects to communicate design research in the way that does not rely on predictive theory. As the authors emphasize that the annotated portfolios are used to communicate design research, I consider them similar to

publications. As for the drifts, however, they are not evident in this format, because the portfolio does not depict intermediate artifacts.

From the examples above, I noted that no drifts are made available to an audience beyond the researchers themselves. A borderline exception is the Critical Journal, which could be shared with design research instructors for educational purposes (Sadokierski, 2019). A thesis is somewhere between a "pure" research paper and an educational outcome, but I could not find directly related information on thesis writing other than described by Krogh and Koskinen (2020). Even so, they focus on the drifts rather than how to document and present the drifts.

Autobiographical Design

Bardzell et al. (2016) consider documentation itself as a research through design process. The method I use to approach designing this thesis is *autobiographical design*. Neustaedter and Sengers (2012) define autobiographical design as: "design research drawing on extensive, genuine usage by those creating or building the system."

Neustaedter and Sengers (2012) arrived at this definition based on interviews with HCI experts and integrated their experience with autobiographical design into a summary of good autobiographical design practices and how learning occurs in autobiographical design.

In order to support genuine usage, autobiographical design requires a functional real system that the researcher uses. Such systems could be basic to start with, and will evolve as the usage grows. The researchers also prefer that the motivation for using real systems be genuine as well, instead of doing research for its own sake. They argue that genuine usage could bring many benefits, such as a deeper understanding of the system in everyday life, and rapid fixing and iterating of the system according to nuanced needs.

Autobiographical design could bring other benefits as well. Autobiographical design could reveal both big effects as well as subtle understandings through the researcher's self-usage. Big effects that are captured early through self-usage could be addressed early, while the subtle understandings could be inspirations for further design. Autobiographical design could also be helpful when there is no existing system of similar functionality, and early explorations would require the researcher to prototype one working system in order to identify the opportunities.

Desjardins and Wakkary (2016) document their project to convert a van into a camper van as an autobiographical design project. Tapping

into the concept of end-user development, the maker community, and home reconfiguration, they were able to distill knowledge relevant to HCI research. The researchers learned from their own genuine experience of living in the camper van and expanded their discussion on situating the design of the Internet of Things in real, lived-in, environments.

Ways to Narrate a Design Research Journey

Outside of standard academic writing of a thesis, which is often analytical and organized by Introduction, Background, Methods, Conclusion, and Discussion, there is a vast landscape of ways to tell a story in the literature world. I take two relevant examples from Jack Kerouac and W. G. Sebald and discuss them here.

Jack Kerouac was an American writer who produced literary works of a myriad of genres, such as fiction, non-fiction, and poetry (Theado, 2000). His widely celebrated book, *On the Road* (Kerouac, 1999), features his transition from a more traditional way of storytelling to what he called "spontaneous prose". Although *One the Road* was famously typed on a continuous scroll of paper in a few weeks in 1951, he started experimenting with this style of writing inspired by literature, religion, and myths through extensive journaling in 1948 (Weinreich, 1990). From *On the Road* and on, his work contained long sentences that resemble an "endless dream" and pieces of memories (Theado, 2000). Instead of separating sentences by periods, he often used dashes that represent moments of breaths. He also combined written language with spoken words as thought associations came to his mind in his sentences. While *On the Road* is strictly chronological, he deviated from this style in his later works, depicting a more dreamy world where sequences of real events took place among the associations.

W. G. Sebald was a German writer known for his "melancholy, ruminative, strange" prose style that produces works hard to be defined as fiction or memoir (Bere, 2002). One notable feature of his books, such as *Austerlitz* (Sebald et al., 2011), is the extensive use of images (such as in Fig. 5). These images are strategically placed in the text, making the readers pause and attempt to decode them and their relationship with the story told (Eshel, 2003). The images are often ambiguous, bringing the readers "into an unreal world unknown to them". They also blur the line between the fictional present moment, as captured in the image, where the protagonists are now missing, and the present moment of the readers reading the text.

its own. Otherwise, all I remember of the denizens of the Nocturama is that several of them had strikingly large eyes, and the fixed, inquiring gaze found in certain painters and philosophers who seek to penetrate the darkness which surrounds us purely by means of looking and thinking.



I believe that my mind also dwelt on the question of whether the electric light was turned on for the creatures in the Nocturama when real night fell and the zoo was closed to the public, so that as day dawned over their topsy-turvy miniature universe they could fall asleep with some degree of reassurance. Over the years, images of the interior of the Nocturama have become confused in my mind with my memories of the *Salle des pas perdus*, as it is called, in Antwerp Centraal Station. If I try to conjure up a picture of that waiting room today I immediately see the Nocturama, and if I think of the Nocturama the waiting room springs to

Figure 5: Part of page 18 of W. G. Sebald's *Austerlitz*, credit Sebald et al. (2011).

The discursive narration, as in the "spontaneous prose" and the discursive use of images by Sebald could be alternative ways to tell the story of design research. For example, Sabald's work was cited as one of the inspirations for a mix-reality design research project that investigate how we may evoke experiences of the past memories associated with a place, similar to the theme of Sebald's writings (Fu et al., 2021). Experimentation with writing can be found in the doctoral thesis, *Interruption: Writing a Dissident Architecture*, by Karami (2018), which contains interweaving formal discussion and analysis with fictional stories of various writing styles and formats. These written forms have informed design research in building artifacts, debriefing, and communicating knowledge.

Journey

When I talk about "thesis" here in this narration, I am referring to the journey of developing the tangible artifacts and their supporting systems, not the bigger thesis you are currently reading. You might find the Chapter, How to Read This Thesis, to be a helpful map in this wild landscape, but a headfirst dive into the Journey could be an adventure as well.

Many of the implementation details are left out in this chapter for the flow of the narration. However, you can find them in Appendix A: Technical Implementation.

The Start

> This thesis started with a pouch with a sticky inner surface, which I developed in Professor Lining Yao's class in the Spring of 2022.

The development of this thesis started with existing artifacts, which were pouches with sticky insides that could temporarily hold different shapes as in Fig.6 and therefore are reconfigurable and reusable. The pouches were developed in Professor Lining Yao's class at Carnegie Mellon University, *Inclusive Tangible and Material Interfaces Spring 2022*, with inspiration and help from Violet Han and Di Wu.



Figure 6: The original sticky pouch, developed with Violet Han and Di Wu in Professor Lining Yao's class in Spring 2022.

Pneumatic Tangible Interaction Experiences There has been an increasing number of works in non-rigid tangible interactions in recent years (Boem & Troiano, 2019). In their 2019 review paper, Boem and Troiano (2019) observed that non-rigid input devices come in many shapes, materials, and sensing methods, and they have a wide range of applications.

One particular form of non-rigid device is the pneumatic device, the structure I am using. Pneumatic devices can be actuated to deform as displays as well as sense external forces or their own deformation (Boem & Troiano, 2019). The compliance of pneumatic devices could offer novel experiences to users.

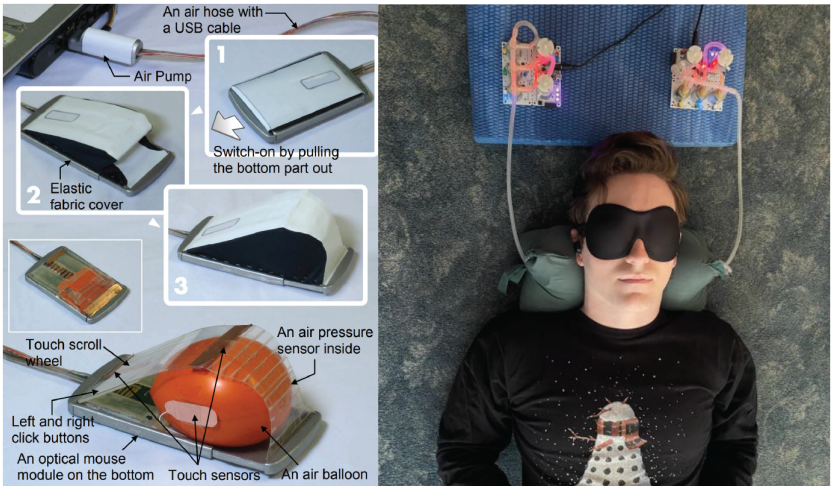


Figure 7: Left: Inflatable Mouse, credit Kim et al. (2008). Right: Neck Hammock, credit: Kidder-Wolff (2022)

With the novelty, many works focus on replicating the functions of standard input devices and adding a little extra functionality. For example, the Inflatable Mouse project (Fig. 7 Left) uses an air balloon both as a sensor for detecting the user's gesture and as a source of passive haptic feedback (Kim et al., 2008). They used pressure sensors and touch sensors. With this device, the user could experience the traditional scrolling, clicking, and zooming while feeling the pressure of air in their hand, rendering a different tactile experience.

There are also works that explore new experiences and functions unique to pneumatic devices. Kidder-Wolff (2022) designed a "neck hammock", a pillow-like device embedded with two air pouches that would inflate and deflate in a pattern that generates a "rocking" sensation when he lay his head in between (Fig. 7 Right). The softness and the yield of the balloons are necessary to create such an experience. The other device, named Inflated Roly-Poly, consists of an inflated PVC balloon with a heavy rounded base like a roly-poly toy (Kwon et al., 2012). The users could press, punch, hit, and push the device to interact with video games. Since the material is light and soft, it creates novelty in interaction methods and makes users more engaged.

> After that class had concluded, I had not wanted to do anything specific with that pouch yet, but was focusing on something completely different for my thesis. However, my summer internship brought me back to the idea of designing tangibles instead of getting lost in the virtual world, so I picked up the pouches once again.
>
> Looking back from where I am now, there were definitely personal motivations that shaped the thesis into what it is now. First of them is my fascination with the human body. From the pouches, I decided, a somewhat forceful stretch, that I wanted to put it on someone's

body. Maybe it was the tactile feeling of a soft material, maybe it was the inflation that was reminiscent of a fish's sac.

Somaesthetic Design This thesis is inspired by the concept of somaesthetic design, which hypothesizes that awareness of the world and potential improvement in the ways we live and act could be achieved through body awareness. (Höök et al., 2015). Somaesthetic design is related to embodied interaction design as they both consider tangibility as important, but somaesthetic design emphasizes designing for experiences (Höök, 2018). It is also different from affective computing where feelings are studied as objects, and instead includes the felt experience in the world. As for aesthetics, somaesthetic experiences are completed in dialogues between the users and the design. Some reading notes on this topic are included in Fig. 8.

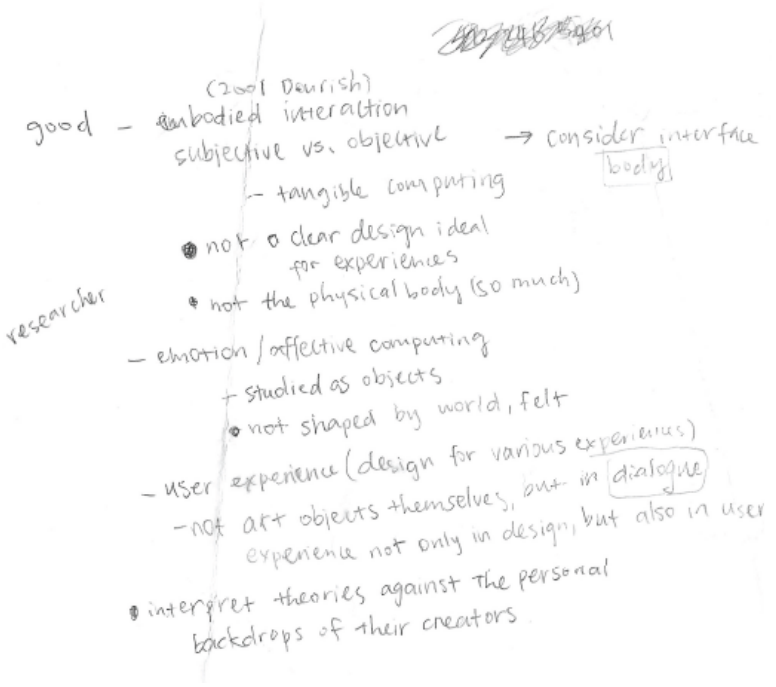


Figure 8: The notes I took from reading the preface of *Designing with the Body: Somaesthetic Interaction Design* by Höök (2018)

Schiphorst (2006) raised the concept of "wearing oneself", which refers to carrying external devices with the wearer's own body-generated data to enhance the awareness of the self. Schiphorst designed a wearable skirt that is networked to communicate the wearer's breathing information to others as a performative experience (Fig. 9 Left). Similarly, Demir et al. (2022) captures her own experience of chronic pain in a wearable piece that could evoke the abstract feeling of pain in its wearer (Fig. 9 Right). Both precedents



Figure 9: Left: *exhale* by Schiphorst (2006). Right: The wearable from Squeaky/Pain, credit Demir et al. (2022).

draw the wearer's attention to the body and leverage this attention to convey their design messages.

Shape-Changing Wearable Interface Soft materials are good candidates for on-body interactions. Clothing, in particular, is a fruitful field of soft technology because it could be both expressive and intimate. Wearable sensors could take in useful information such as body gestures and wearable actuators could be aesthetically pleasing and functionally useful as well. PolySense is a project that adds piezoresistive properties to textiles via in-situ polymerization so that the fabric could sense pressure and deformation (Honnet et al., 2020). ClothTiles are made with 3D printed soft materials and shape-changing alloy to enable actuation (Muthukumarana et al., 2021). These two examples find that the body is a rich source of information and display.

> The other line of fascination was simply my ambition and fantasy to make something useless yet provocative, and I find myself situated in speculative design.

Speculative and Discursive Design This thesis operates in the framework of Speculative Design. In *Speculative Everything*, Dunne and Raby (2013) proposed that Speculative Design holds the belief that by considering alternative futures, we could better understand the

present and open up opportunities to move forward to a preferable future.

Speculative Design is the opposite of trying to pinpoint and predict the future (Dunne & Raby, 2013). It operates at the intersection of what is probable and what is plausible. The probable futures represent ones when the current world trajectory does not change radically or will not be disrupted by disasters. The plausible futures propose alternatives and involve "scenario planning and foresight". Speculative Design hence tries to bring what is plausible into the probable, to prompt people to actively discuss and participate in building what they prefer. A closely related concept is discursive design. Designed to foster understanding and conversations, this thesis falls into the category of discursive design as well. Instead of developing tools that serve specific purposes, discursive design seeks to convey ideas and ideologies around social and psychological topics (Tharp & Tharp, 2018).

Design with Ambiguity W. W. Gaver et al. (2003) identify the value of incorporating ambiguity into design practices and classify ambiguity into three main categories: the ambiguity of information, the ambiguity of context, and the ambiguity of relationship. The ambiguity of information arises when the audience is not sure of the information presented. The context is ambiguous when the artifact could exhibit different meanings in different settings. The design has an ambiguous relationship with the audience when it prompts the audience to reflect on how they may inhabit, use, or simply approach the design.

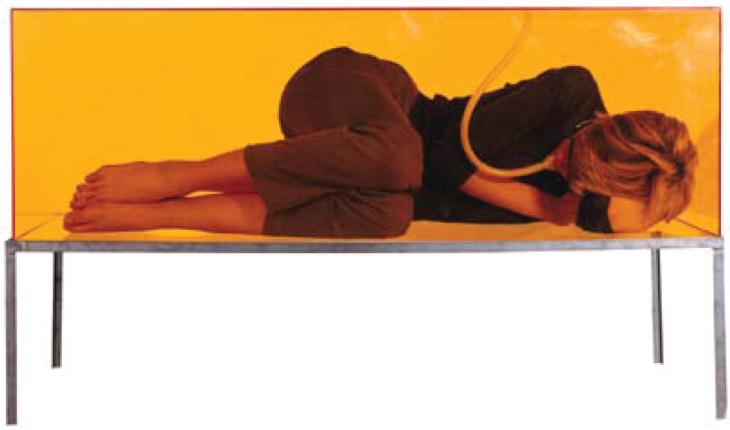


Figure 10: Faraday Chair by Dunne & Raby (1995). Photo credit: Lubna Hammoud

One example is the Faraday Chair by Dunne & Raby (1995), which

is at the same time protective due to the shielding of electromagnetic waves but also restrictive due to its hard plastic closure (Fig. 10). Another example is the Home Health Monitor by Boucher et al. (2001). The project mapped sensor readings from home to sentences from a database of horoscopes and generates a sense of ambiguous "fortune-telling" language leveraging the inherent vagueness in the horoscopes. The sensors installed in homes do not capture everything happening in the homes and thus are not a reliable source of information, yet they mean something, just as the horoscopes always get something right.

> These two lines faded in and out of the development as I will discuss later. The third concept in my mind was the concept of intimacy. I did not think too much about why, because it seemed to be a natural craving for my situation as a graduate student, and the human body already hinted at the concept of intimacy. Later in the development, I also favored making life harder for people to do simple tasks, thus the concept of effort, and having a sophisticated outcome as the side effect of the simple task.
>
> While personal interests shaped the earlier inquiry, and only subtly influenced the later decisions, implementation motivated another major portion of drifts throughout the thesis. I had to admit, that the reason I convinced myself to pursue speculative design was partly that the pouches I implemented were not precise enough to be used as a scientific instrument or sensors. Ambiguity came into play for this reason as well. I first characterized the performance of capacitive sensing based on the Multi-Touch Kit work,

Sensing Implementation

L. Yao suggested that I consider capacitive sensing as a way to digitally capture the shape transformations of the pouches. With Professor Daragh Byrne's suggestion of manufacturing custom capacitive sensing structures with the Multi-Touch Kit (MTK) (Pourjafarian et al., 2019), I downloaded the code and built the structure as they specified.

> from which I learned that the Arduino [1] analogRead signal was linearly proportional to the capacitance of the sensing electrodes I attached it to.

Next, I measured the signals from a single MTK electrode pair and came to the conclusion that the analog signals from the Arduino were roughly proportional to the capacitance between the electrode pair.

[1] I used the Elegoo Uno R3 Board, https://www.elegoo.com/products/elegoo-uno-r3-board

> I went on to build what I referred to as a slider pouch (Fig. 11), a pouch with a tapered sensor that records roughly where you pressed into the pouch and made the sticky layer stuck to the bottom.

With this observation, I made a slider pouch, which was a two-layer pouch with a tapered triangle electrode on one side and a rect-

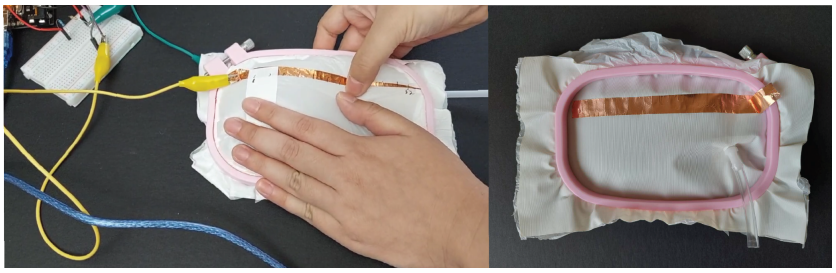


Figure 11: Left: the slider pouch in action. Right: the back of the pouch, with a uniform rectangular electrode.

angular electrode on the other side. These electrodes were all made out of thin copper tape. The inside of the pouch was completely covered by one sheet of sticky tape, so I could pinch anywhere along the electrode and it would hold the contact.

> It was through this pouch that I learned about how unreliable the sensing method was. The contact area, which determined the signal strength, could be changed with either the location of the contact, i.e. how far you pressed from the tapered end, or the area of the contact.

By varying the distance from one end to where the electrodes are in contact, the width of the contact changes, and thus the capacitance and the signal.

I wrote a web interface to visualize the distance from one end of the electrode to my finger. With the visualizer, it was easier to see that the read signal did change with the distance, but it also changed with the length of the contact area, since I pinched on the electrode with varying force.

> This made the signal difficult to interpret.

After the slider pouch, I constructed the simpler version of the MTK, which detects the contact between crossing electrodes. I made two horizontal lines and two vertical lines with copper tape (Fig. 12), and printed the MTK values to Arduino Serial [2] (Fig. 13). Through the Serial Plotter in the IDE, I could see the value changed with the pattern I touched the electrodes and with where they were stuck to each other.

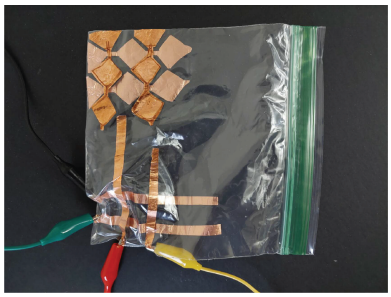


Figure 12: The simpler pouch with crossing electrodes.

I switched to pouches with sticky patches inside instead of using a large sheet of tape, such that only the crossing part of the electrodes could stick to each other and be peeled apart with ease.

During this encounter, I also implemented the Python [3] interface to the Arduino Serial, and collected the Arduino data into usable queues, ready for further processing.

[2] Serial Communication is explained here: https://en.wikipedia.org/wiki/Serial_communication

[3] Python 3.10.4 | packaged by conda-forge, https://www.python.org

> You may ask me why I needed to sense the contact location. As inherited from the fascination with bodies, I had wanted to make it something tangible that could serve as alternative inputs to regular computing systems. This is what I think now I was thinking back then, though the

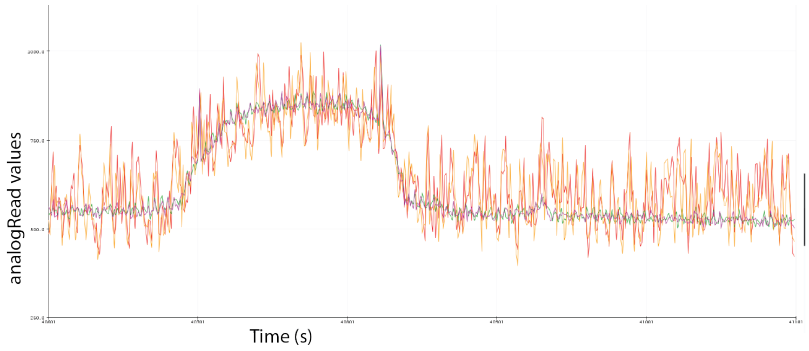


Figure 13: The Arduino analogRead response to a pinch on the electrodes.

argument was very ill-defined at that stage in the Fall of 2022. (In the forecast, I wrote the imaginary evaluation and contribution which never came to light for this encounter.)

I had the habit of talking extensively about my thesis with people beyond my advising team, such as my friends and my mother. When I was grappling with the ambiguity in the pouch signal,

The characterization of the MTK signals and the design of the slider pouch were the attempts to directly infer the pouch states based on its signals. However, after I manufactured the slider pouch, I realized that the signals from the pouches are hard to interpret and quantify. With a human eye, I could see the values go up and down, as in the figure above, but the absolute values would change with small variations in the setup. The changes were also gradual, with rising and dropping windows.

it was my mother (Yanli Xie) who suggested that this could be a design opportunity, since it was a big contrast with common digital devices.

I noted that sensing through pouches in this manner creates a sharp contrast with sensing with buttons, which are nearly binary and much more stable.



I took note of this idea and formally introduced the concept of ambiguity.

Combining the above observation with the intention to design an on-body system, I decided that the uncertainty in sensing would be

a good mapping of body gestures to some form of output because body gestures were ambiguous in their meanings as well.

> My other friends helped with brainstorming what kind of expressive form this ambiguity could be mapped to. Jingyang Liu and Sinan Goral both mentioned light and sound,

I received several suggestions from Sinan Goral and Jingyang Liu that encouraged me to consider light and sound as outputs.

> but to me, these two modalities were too implicit and decorative to be provocative.

However, I considered these modes of outputs to be not explicit enough, because they could be interpreted as purely performative, and the inherent mapping from movements to intimacy could be obscured.

> The struggle with output modalities has proved to be a constant theme throughout my drifts in the thesis. With the concept of ambiguity, I drew inspiration from an article on social media that described several ambiguous "lover's talks", or something similar. The only one that I can remember and translate to English was "I've never seen anyone like you", which could be a compliment for uniqueness or a judgment of a weirdo.

In the meantime, I came across a sentence, "I've never seen anyone like you", which could be a compliment or a pejorative judgment.

> I was deeply impressed by the ability of language to be both explicit and ambiguous, yet found it hard to argue what made language ambiguous in the first place.

This form of ambiguity struck me, and I liked the seemingly explicit yet elusive quality of language.

> In one review session, Professor Marti Louw brought up Edward Lear's limericks, which probably defined ambiguity along the line of having multiple interpretations or, not having a sensible interpretation among them. I was particularly moved by the *There was a Young Lady of Portugal* (Fig. 14).

I was also pointed to Edward Lear's limericks by an external reviewer Professor Marti Louw. Lear wrote a large collection of short, rhymed limericks the meaning of which could be greatly extended with imagination.

> Feeling quite confident about my skills on natural language generation, I aspired to make something like *a Lear's nonsense generator*.

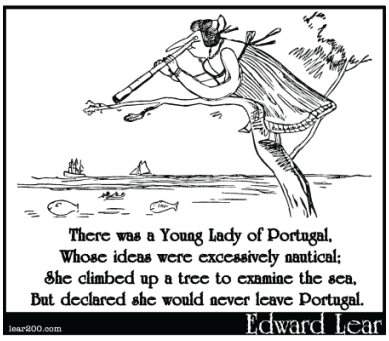


Figure 14: Edward Lear's limerick: There was a Young Lady of Portugal.

Generating Ambiguous Language

Inspired by the literature review, I dived into natural language generation. I was aware that I would not be able to train a language model from scratch, so I hoped to utilize the existing models and build on top of them. Specifically, I was hoping to map body gestures in space

to word embeddings in their vector space. In order to generate language about intimacy, I brainstormed the following lists of words, as in Fig. 15. The first is a list of verbs, which could be used to connect the nouns in the second list and be modified by the adverbs in the third list.

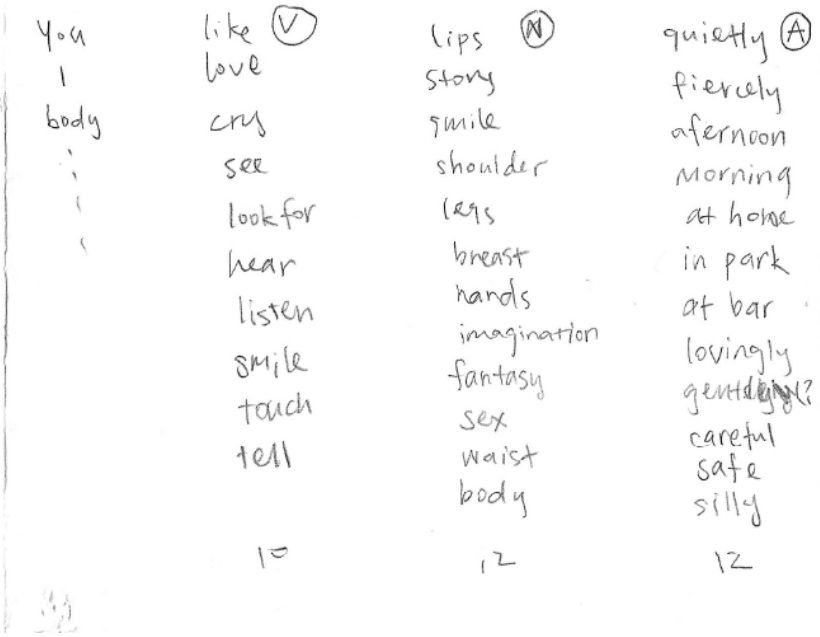


Figure 15: Three lists of words I thought about using to generate ambiguous language.

As for how intimacy came into play, I had multiple vague ideas back then. I had wanted to put it near the joints, for example, so that it could capture the wearer's body movements and gestures. It was also uncertain who the wearer was being intimate with. Sometimes I thought it could be a two-person play, and the intimacy would play out like many HCI papers fostering physical interactions that I had seen. Sometimes I thought about solitude, and how a person may hear their own voice talking to them about their most intimate partner: the body they reside in. This concept evolved briefly into some ideation on body image, and how maybe the pouches on the wearer's body could trigger, in the form of audio playback to the wearer, different thoughts and judgments one may have on their own body.

This idea was quickly abandoned because I wanted to engage more with the material, not the message itself, and this topic was too "intimate" to me. Nevertheless, I leaned into talking about intimacy with oneself, and hypothesized that the pouches could capture intimate gestures, map them to intimate language, and influence how the wearer feels.

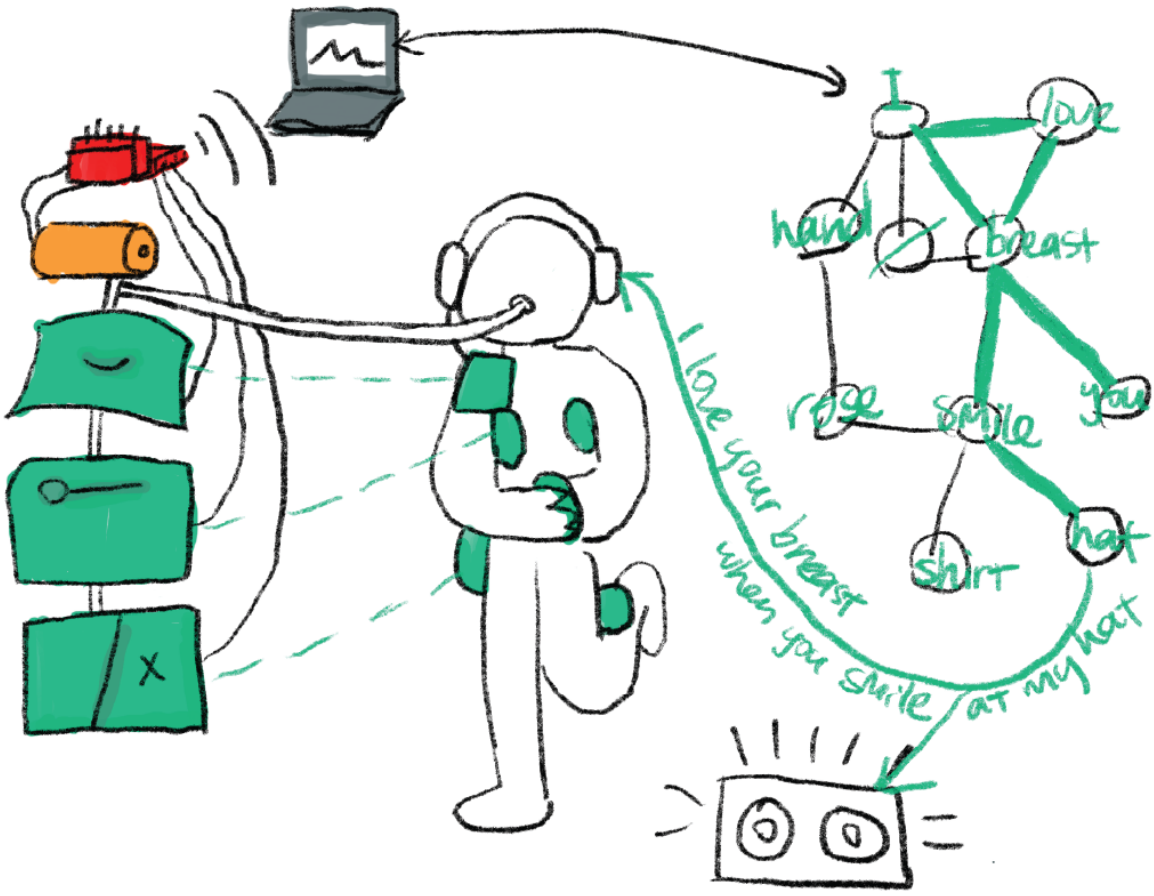


Figure 16: The components of the proposed system: mechanized on–body pouches with sensors feeding into a computer that generates related ambiguous language to the wearer.

Landing: Intimate Ambiguity

We perceive and judge our bodies constantly when we are moving through the world as well as when we are alone in front of mirrors. Sometimes we like our bodies, and sometimes we do not. Nevertheless, the body is the most intimate environment we all live in, and it is valuable to explore this environment and reach a homeostasis within it so that we could reach outwards beyond its constraints. Language, on the other hand, is an outward projection of ourselves, yet it could have profound effects when it is presented to us. This thesis captures the ambiguity in one's attitude towards oneself through shape-changing materials and verbalizes such quality through equally ambiguous language about self-intimacy, with the intention to prompt the wearer to clarify for themselves what they feel.

More specifically, my thesis raises the following questions:

- How can I leverage the shape-changing potential of air pouches to design continuous input devices that map uncertain body gestures to ambiguous language about intimacy with oneself?
- What effect would the language about intimacy generated by gesture have on the wearer around body perception?

This thesis addresses these questions by developing and testing an on-body system of air pouches that serves as sensors of body gestures, and an associated natural language model to generate language about intimacy. Each set of similar body gestures is mapped to elements in the language model, which is customarily trained to produce ambiguous output in the form of poetry.

The main contribution of this thesis lies in the following areas:

- An on-body novel interactive structure that senses body gestures in the form of continuous output.
- A custom language model that outputs ambiguous language based on signals from tangible objects.
- An interactive experience of self-intimacy that opens the conversation around what wearable digital devices could engage discussions around one's body.

> If things went well, I would not be writing these very words. Several things happened during my work on intimate ambiguity. First, I did not find any literature on generating ambiguous language. Ambiguity, as its name suggests, was ambiguous, and it defies all kinds of measurements. I was overly confident in my ability to tackle natural language generation tasks, and this task might as well be another entirely different thesis.

I discovered two problems with the approach of piecing words from the three lists together: it was unintuitive where the word embeddings will be, simply by looking at the lists; it was also hard to perceive the randomly generated sentences as ambiguous rather than plainly random. The machine's interpretation of the words only captures what they mean on average and the machine does not have a quantitative measure of ambiguity.

In fact, to bring tangibility back, I had a brief proposition on collaborative poetry generation, in which a couple of participants would blow into the pouches and manipulate them at the same time, as illustrated in Fig. 17.



This tangent was also abandoned for the practical reasons above. Second, when I presented the argument to people, especially S. Goral, it was very difficult to get them to conceptualize the ambiguity in relation to body and language in the same way these words made sense to me. I had three different definitions, "ambiguity", "continuity", and "uncertainty", each meaning something slightly different, which I would spare you from understanding.

Continuity of signal = signals that can change at arbitrarily small steps

Uncertainty of action = having varied and unpredictable responses from similar actions

Ambiguity of information = capable of eliciting multiple interpretations from the receiver of information

Third, I could not find a mapping from intimacy to language. I was going to bridge them with emotions and perceptions about self.

Mapping Body Gestures to Intimacy

To ground the language generation, I needed the part that mapped body gestures to factors about intimacy that would create meanings in the ambiguous language. Here, I found my core assumption that the body gestures pouches could sense would map to emotions (Fig. 18). I already argued that this mapping would be somewhat uncertain, yet it needed to exist to build the algorithm.

I created a mindmap that had emotion words at the center and body parts around them, as in Fig. 19. I reflected on my habits of



Figure 17: A brief digression to collaborative poetry.

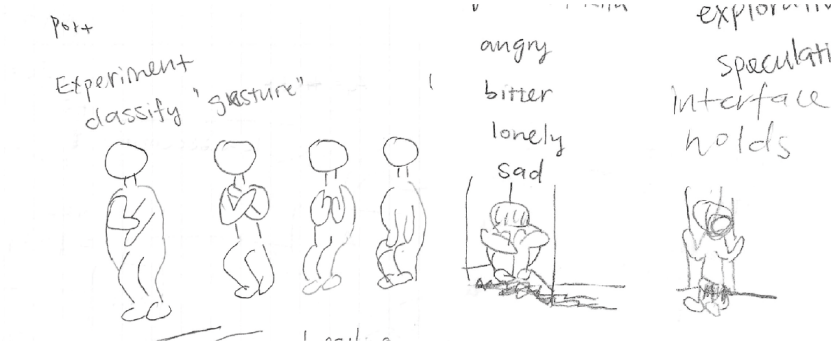


Figure 18: Attempts to draw a few gestures.

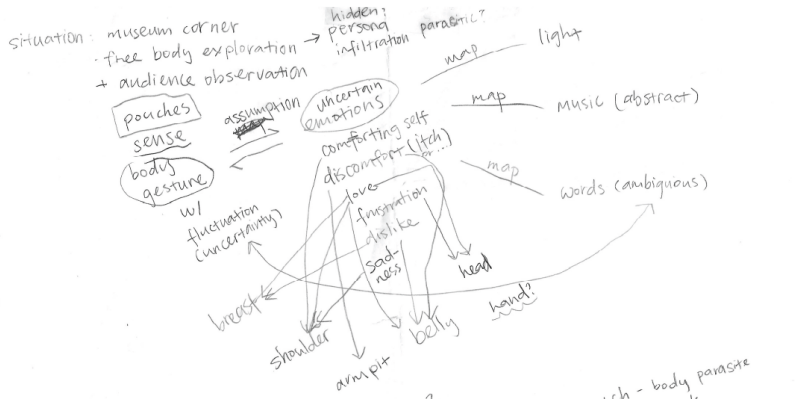


Figure 19: Some thought and emotions are mapped to different body parts and some thoughts on output modalities.

touching my own body, and connected what I felt at that moment to the body parts I would be touching.

> For example, when I pat my shoulder, I might feel sad, so the system could say something strange like, "I love your shoulder." as in Fig. 20.

For example, when I put my hand on my shoulder, I could be comforting myself, showing love to myself, or expressing sadness. There are other more ambiguous body parts, such as the breasts. When I touch my breast, I could be expressing self-love, and I could also be expressing negative judgment.

> After a deeper dive into what I would do to my body when I felt different emotions towards myself, I did not find explicit mappings that are meaningful enough to be hard-coded into the mapping function.

How could I map these two extremes into the ambiguous language that had limited access to the high degree of ambiguity these two emotions required?

Discussion: Intimate Ambiguity

From the preliminary development of the system, I identified a design opportunity in the continuous and unstable signal from the capacitive sensing of the pouches. Starting with the shape-sensing uncertainty and my personal interest in wearable systems, I hypothesized that the signals from the pouches would capture the wearer's body gesture with ambiguity. Furthermore, I hypothesized that body gestures would map to perception and attitude towards self, with ambiguity as well. Lastly, this chain of ambiguity would be encoded in poetry as ambiguous language. I would like the wearer to reflect on the poetry and subsequently their perception of and emotions around their body.

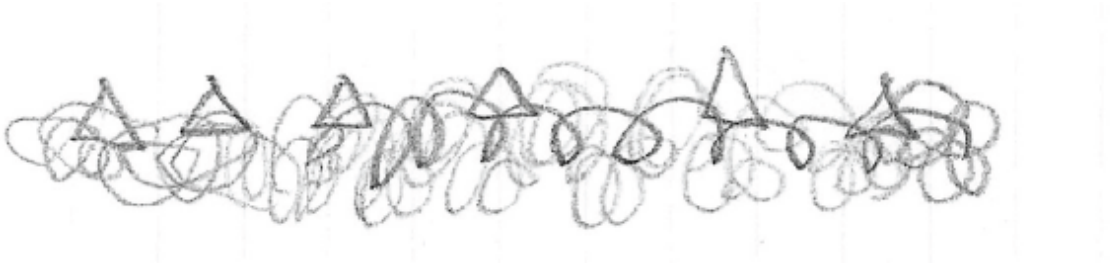


Figure 20: Fantasy: "I love your shoulder."

However, the development of the natural language generation model as well as the body-emotion map revealed parts of this thesis that I did not anticipate during the framing process. In short, it would be difficult to generate controlled ambiguity with the existing language models, and the chain of ambiguity above is too ambiguous to be perceived as not random. I also received feedback from S. Goral that the definition of ambiguity and its related words are confusing, probably inherently so.

These unexpected observations brought several of my hidden assumptions to light. I assumed that machine learning algorithms could capture "ambiguity", however, this concept was ill-defined in my context. There are researchers working on aesthetic evaluation

of AI-generated poems with methods in psychology and statistics (Hitsuwari et al., 2023), yet those works would lead me further away from the wearable artifact itself. When I claim this thesis to be a discursive design project, it loses the way traditional engineering might be applied to solve a problem and became "under-constrained" as Zimmerman et al. (2007) suggested. I realized that I was wrestling with unbounded possibilities to interpret my artifact, and more work was needed to convey an idea beyond what would appear as random.



In summary, during this encounter, I proposed a plan to design the concept of ambiguity with continuous and unstable capacitive sensing signals from a novel material structure. This structure was designed to bring awareness to the user's body and encourage them to explore body perception and emotions towards the body. Technical successes were achieved in manufacturing the multi-touch pouches and interfacing with Arduino and Python. Through the development of the system, however, I learned that ambiguity was very hard to quantify and reproduce, and the mapping from an individual's body gesture to their emotions was also an ill-defined problem. I had only designed the system based on my observations, and I was aware that if I had more samples or more data, the above problems might be able to be solved.

> At that stage, I was ready to abandon the concept of ambiguity, for its ambiguity. Following the conversation with S. Goral, which led me to the thoughts of collaborative poetry and its various versions of human-machine antagonistic relationships, I started to implement alternative sensing mechanisms for the pouch shapes, as well as algorithms to control the pumps that would bring the system its own agency.

Since the mapping from emotion to gestures and the natural language generation for ambiguity was too open to tackle, I made the above mind map to think through what else I could do with machine learning (ML). Fig. 21 classified ML in terms of its input and output and arrived at the Explainable Artificial Intelligence (AI) and Tangible AI concepts.

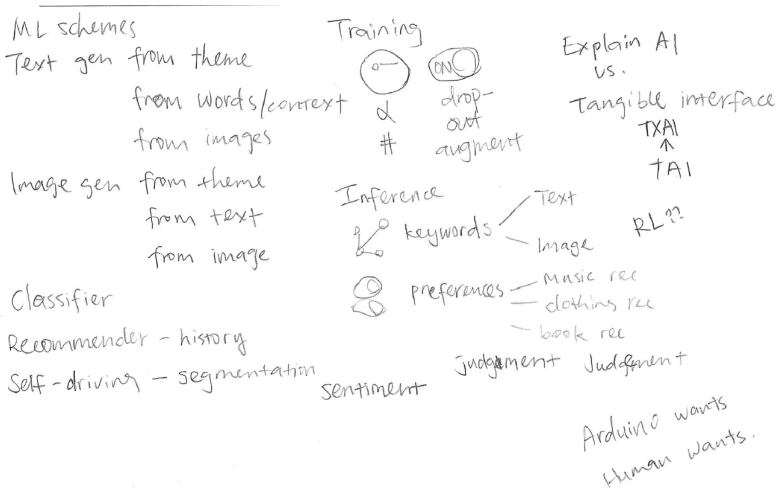


Figure 21: Collection of ML schemes, with different input and output modes.

> What I did not anticipate was that the sensing mechanism eventually revived the concept of ambiguity, or uncertainty, or any of that fuzzy ball of concepts at later encounters of the development.
>
> Now when I look back to when I started coding finally, I realize that I already had the tendency to throw feasibility and the material world out of the window. First of all, I wanted to create an intelligent pump control system that would play with or against the human player of the pouches by inflation and deflation of desired amounts.

Control Implementation

I started by implementing a simulation system such that I could test various of the ML schemes without having to manufacture all the hardware and running the risk of bursting many pouches with poorly programmed pumps.

The first system for simulating pouches was simply a value that can be changed at different rates linearly proportional to the value of the simulated pump speed. This modeled a sealed pouch and an ideal pump.

> Without much thinking, I programmed an ML regression model that calculated the pump speed, which related to the volume change of the pouch by a constant factor, and the model happily calculated for me how fast to run the pumps in order to keep the pouch volume in desired ranges.

My first ML model was a regressor model. It took a sequence of ups and downs as "button states", which corresponded to whether the pouch was popped up (no contact) or pressed down, respectively. The output was the required pump speed at each time step in or-

der for the human operator to perform the action. For example, if the pouch was previously pressed down, and the human operator wanted it to pop up, the pump would need to pump more air in to inflate the pouch and pop the area that was stuck.

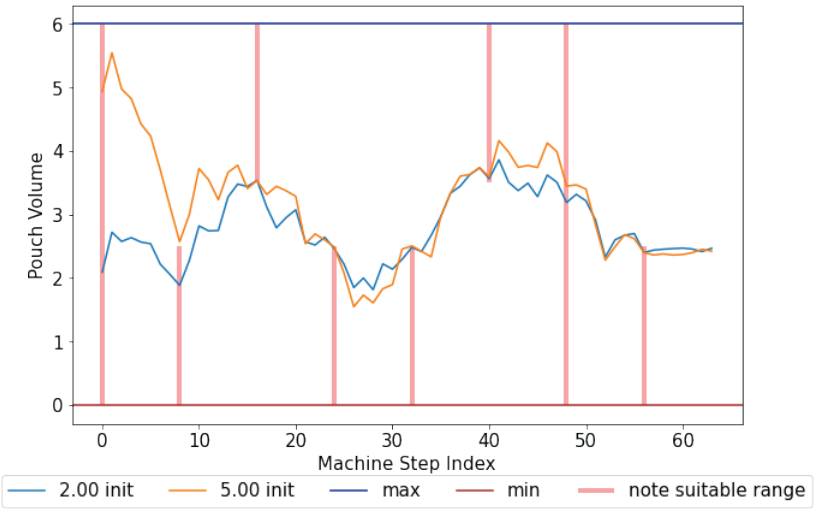


Figure 22: Output of a simulated pump with ML regression of pump speed. The orange and blue curves are pouch volume curves as the results of the pump speed.

However, I realized quickly that in reality, the linear factor does not exist, and the regression model was too simple, because my pouches could be leaky, the pressure change would change how the pumps worked, and human actions would also be variables of the system.

This regressor worked pretty well given its inputs and outputs as in Fig. 22, but it turned out to be not very useful or realistic. Given a target volume for the action to happen, the required pump speed was simply the volume adjustment divided by time, which did not need a neural network to run. The realistic scenario was that the pouch might not be completely sealed, and the pump inflation rate might change with respect to the air pressure inside the pouch as well.

Soon after that attempt, J. Liu suggested that I look into PID control, but I thought it would be too limited and not as intelligent as Reinforcement Learning (RL) methods, which could drive cars on the highway. All this mess was comparable to highway traffic, I thought.

I started to implement an OpenAI Gym[4] environment for Reinforcement Learning (RL), because I thought an intelligent agent operating the pumps could take into account the complexities in the dynamics more directly. The input was the same sequence of "states" and the output would be pump speed at each timestep as well.

[4] OpanAI Gym 0.21.0, https://github.com/openai/gym

Here I had a huge tangent into writing a good simulation of realistic pouches. I was going to flex my shaky muscles on the thermodynamics equation I learned in my undergraduate program, $PV = nRT$, and wrote a pouch simulation that took into account all of the factors

like pressure, volume, and amount of air molecules.

I started by listing what humans and machines could do, such as pressing, deflation, and inflation, as in Fig. 23. Then, for each of the

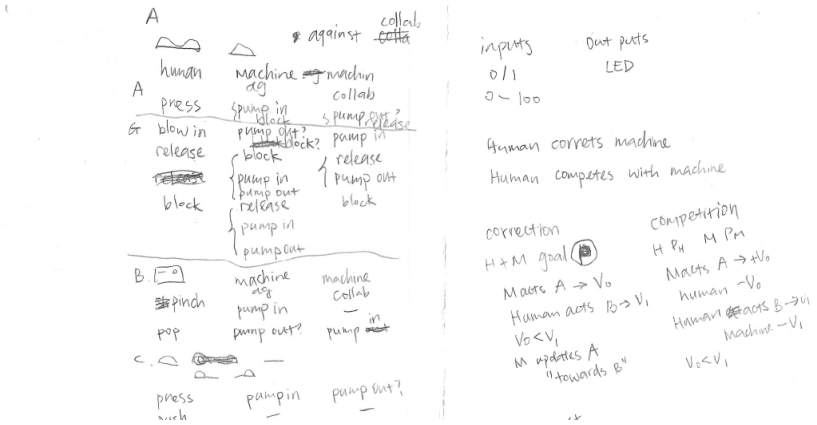


Figure 23: Actions and effects the human operator and the machine pump could take.

actions, I modeled the pressure, volume, and amount of air in the pouches as the result of the actions via the equation $PV = nRT$. I noticed a lot of naive assumptions when I first coded it and ended up modeling the full PV curves (Fig. 24).

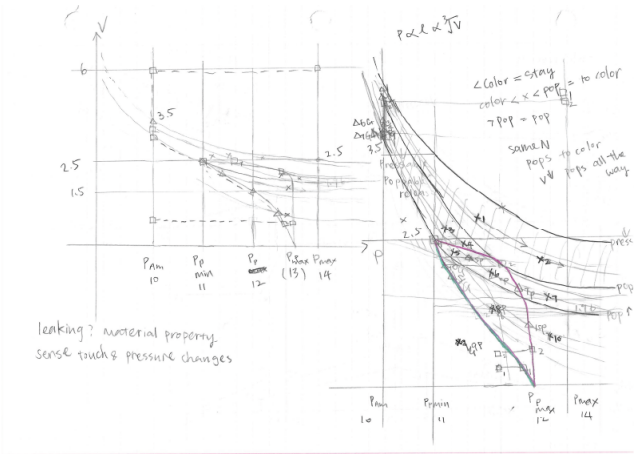


Figure 24: Hand-calculated PV curves and the transitions between different pouch states.

I did nearly write the simulation, but on paper with squares and triangles and rough curves of inverse relations. When I was ready to write the simulation in code, I made the right decision to revert back to test RL on what I thought was a minimum working simulation of an air chamber, maybe. Then I failed gracefully on this simulation, which I reduced and reduced to wildly unrealistic conditions again, not far from the ML regression model.

However, I also had to simplify the system to barely more than the regressor simulation in order for RL to almost work, and it worked

with much more information exposed to it than what would be physically realistic. For example, the agent had to know the current exact volume to decide the next step.

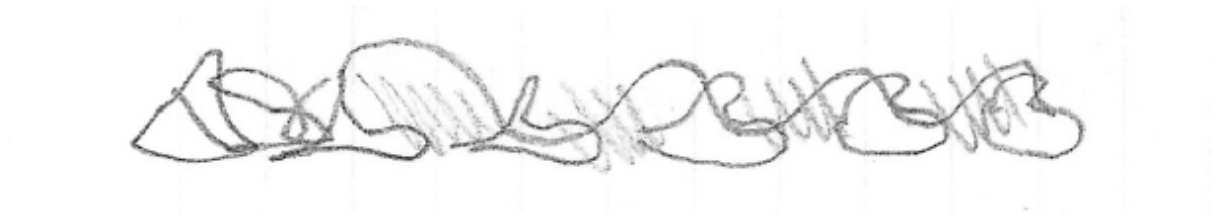
During the development of the RL agent, I realized that extensive training was needed even in the simplest simulated environment.

Music Implementation

> On my way back from CVS one day, I was humming mindlessly some tunes. It was not anything original, and I did not remember anything I sang, but it occurred to me that I could get away with a simpler ML model that only classified shape signals, mapped them to musical notes, and, somehow similar to the MNIST dataset, asked people to label their own notes with gestures.

The difficulties with training an RL agent to control pumps in the real world prompted me to replace active controlling with passive interactions. The passive interactions are one-way: the algorithm tried to detect pouch shapes based on capacitive sensing signals, as in the previous encounter of this thesis.

> Once trained, the user would be able to repeat the gestures and the machine would be able to recall the associated musical note. People would have to sing, I reasoned, because they would be manipulating the pouches with hands, and singing would be easier for me to test as well when it did not obstruct other senses I would be using to code and debug.



As for the output of the system, I opted for sound generation in the form of music performance. This mode of interaction is hand-free, which allowed me to interact with the pouches with one hand, control the code with the other, and sing to the prototype while training. Besides the practical reason, I considered it a valuable property because music could help people get fully immersed in the interaction with different senses.

Sound-based Interface The artifact of this encounter is designed to be void of graphic displays to fully engage the wearer's retrospection. Johansen et al. (2022) review the use of real-world sounds for soundscape compositions. They discuss that soundscapes could engage

not only hearing, but also the immersive experience of all the sounds and their environment. In other words, the context could be "felt" through hearing.

> This was a remarkable realization, which defined the later algorithmic pipelines in the rest of the thesis. It would be difficult to create such a labeled dataset, I hypothesized, which fitted my obsessions with asking people to put in the effort,

Design with Effort This encounter in the thesis increases the complexity of actions involved to achieve an otherwise simple task of recording. Literature has shown that physical and emotional effort could be a powerful design tool in persuasive games (Lyons et al., 2012). Tuned to a proper amount, the effort could strengthen the massage the designers wish to convey.

Calibre is an artifact by Jang (2013) that is an unusually heavy pen held by a stone base (Fig. 25). It is intended to remind its users of the importance of what they are writing. Such users have been world leaders, who the designer hope could consider more what they write may affect the world. By making the pen harder to use, Jang brings attention to the action. Similarly, this encounter of the thesis makes recording harder to encourage contemplating about what is needed to be stored in external memory.



Figure 25: Calibre, credit Jang (2013)

> and it could be expressive, as in the musical notes people sang. I should have noted here that musical performance was something I truly had fear of, which many people later confirmed, and proceeded with caution. But I did not.
>
> On the side, I had read about Agre's critical technical practice (Agre, 2014) and Lucy Suchman's argument about how end-user should be included in the design of technology by then (Suchman, 2002). I became passionate about how we should bring the revolution of AI closer to people using them.

Reflections on AI Agre (2014) introduced the concept of Critical Technical Practice in his research relating to AI. He observed that in the field of AI, researchers use their own systems of language to describe both technical and mundane activities. This use of language eliminates the way to think at a distance from their practice and makes AI a closed field that can only perpetuate its existing structure.

Similarly, people outside of the AI research field could be confused by the mixed usage of formal and informal terms as they mask the complexity of the computing systems. Standing at a distance from the field, people, the actual users of such technologies, are sometimes forced to accept the "authoritative knowledge" (Suchman, 2002) of AI systems and developers, only which is considered conventionally

legitimate. This knowledge could inherit biases and ignorance from AI researchers. In practice, the users of AI systems therefore need to rework the systems to fit their own needs.

Suchman (2002) considers the local reworking of technological systems a form of their generation. Kolkman (2015) takes this reworking one step further by suggesting that people could choose to take ownership of technology from the corporate and government. In his project, OpenSurgery, he explores the possibility of building DIY surgical robots as an alternative to the costly service in the US healthcare system (Fig. 26). Advocating for the agency of citizens and questioning the current medical practice, he uses this speculative example to spark conversations. This encounter of the thesis is similar in the sense that I would like participants to experience training an AI system from scratch that listens to their individual patterns, in opposition to having centralized technology companies provide a black box trained on aggregated yet decontextualized data.



Figure 26: Open Surgery, credit Kolkman (2015)

> There was a cluster of arguments I had a lot of trouble telling apart: If I asked people to train their own ML classifiers, and they fail, would they become pessimistic or even anti-ML when only corporate or government efforts could afford training a useful ML product?

Human Labor Machine learning models often require large amounts of data to train properly(Roh et al., 2021). According to a review paper that provides a list of current data collection methods, both data acquisition and data labeling could involve a large amount of manual labor, such as crowdsourcing data and manual labeling. This manual labor is often concealed once the systems have been trained and proved successful.

> Would my system that revealed the fundamental elements of ML demystify the technology and help democratize them? When people succeed in training their own intimate ML models unique to their own voice and touch gestures, is it better than the generic model, and is it more ethical since it does not exploit anybody's labor?

Tangible AI Researchers consider that people's understanding of AI systems could benefit from having tangible experience interacting with them (Ghajargar et al., 2021). Besides understanding, users could better trust and improve the systems both in their performance and the social issues they involve.

AI systems themselves could be a source of tangibility. By adding a physical "gesture" to the AI system before they carry out the action on a shared control, researchers find that the participants felt more collaborative with the machine (Zuckerman et al., 2022).

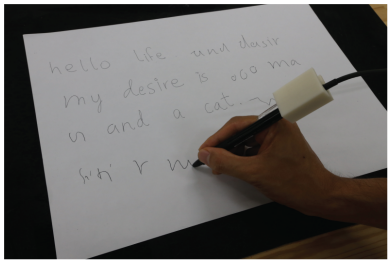


Figure 27: (author)rise, credit Agrawal et al. (2018)

AI could directly act on the human body to create a visceral experience of the technology. *(author)rise* is a project that makes a natural language generation system push the participant's hand to write what it predicts to be the next words (Agrawal et al., 2018) (Fig. 27). Having AI write philosophical words further challenges how people often fail to think about AI's prevalence both in physical and intellectual environments. This encounter seeks similar bodily reactions from the wearer of my system as a result of AI's prediction, but more in the training phase than the application of a trained system.

> Or is the intimate model itself a concept interesting enough that I could open up the space for interpretation to the users?

Non-Rigid Interface for Creative Expression Music is a common medium for expression as the output of non-rigid devices. Troiano et al. (2015) identified 10 works involving deformable materials that serve as MIDI controllers, sound controllers, data sonification, and auditory feedback. For example, the Music Ball Project uses a stuffed embroidered ball to replace the traditional interface of electronic musical instruments and allows novice players to experiment with pre-composed music easily (Weinberg et al., 2000). Troiano et al. (2015) also discovered that musicians were able to adaptively use novel deformable devices to perform, based on custom mappings of gestures and shapes to music parameters. This encounter of the thesis adds another layer to the passive mapping of gesture to music in that I utilize the state of the device as well, which supplements the gestures of the user with its own agency.

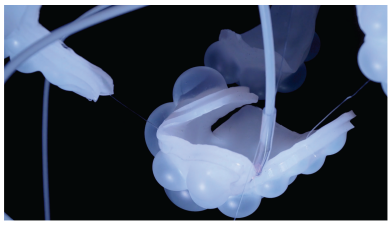


Figure 28: Pero sans Cimon, credit Bourgault et al. (2020)

The form of the soft robotic body could be provocative and expressive in itself. The softness of the material brings more freedom and instability to the artistic expression, just as it brings novel interaction in engineering (Bourgault et al., 2020). *Pero sans Cimon* is a sculpture made of silicone chambers that symbolizes the exploitation of women as reproductive machines in the current society (Fig. 28). Linking robotic softness to machine art practice, the artist-researchers identified that soft robotic materials could be reminiscent of the human body in its curved and bouncy form as well as its pulsing and breathing movements. By placing the soft materials on the body, this encounter of the thesis could further connect inorganic robotic structures to the human body and blur the boundaries of machine art and engineering.



Figure 29: Created by me and DALL-E (https://openai.com/product/dall-e-2).

Landing: Hidden Labor in ML

What if your smart voice assistant has your own voice? What if you have to teach it to speak from scratch, with your own voice, by perhaps reading a 500-page book into its microphone?

Such scenarios could be hard to grasp. The first question taps into the uncanniness people feel when inanimate objects are overly similar to themselves or have intimately close relations with them. The second question reveals more or less the reality of large-scale training of Artificial Intelligence (AI) models, only that such hard labor the large-scale training requires is silently divided and distributed by the organizations that develop them and fulfilled by anonymous groups of people. As AI is entering many aspects of life, a critical discussion around the extent to which an individual could possibly perceive the mechanism of AI and own the AI they use could raise awareness of the deep implications of this technology.



The overarching theme this thesis aims at exploring is how people may feel intimately connected to and directly experience the process of generating an intelligent machine. More specifically, I ask the following questions:

- What *bodily sensations* does the on-body training create?
- What kind of *mental and physical labor* imitates the large-scale training of machine learning models?
- What differences in *feelings, attachments*, etc., are there between a generalized model and an intimate model unique to you?

My thesis provides insights into these questions through the development and evaluation of InVest, a wearable system that combines real-time machine learning with creative music generation. The four components of InVest, the ML model, the gesture sensors, the music mechanism, and the wearable structures, act together to spark reflections and conversations around a speculative AI development experiment that is extremely personalized and intimate.

The main contributions of this thesis lie in the following three areas:

- A wearable system for interactive machine learning with haptic and audio feedback.
- A framework for materializing labor in machine learning with an emphasis on individual's mental and physical labor.
- An intimate embodied experience of machine learning- facilitated music creation.

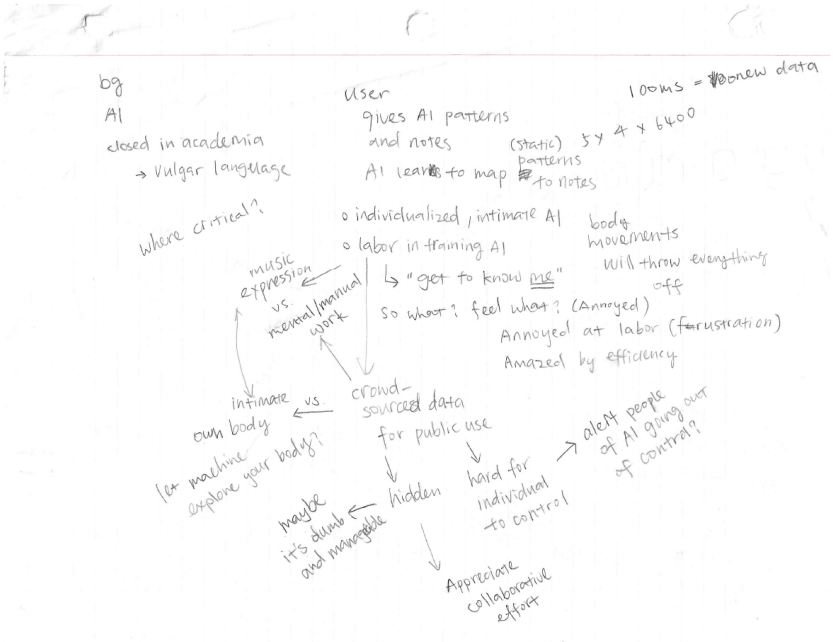


Figure 30: Brainstorming for concepts related to ML and labor.

With my head swirling in these tangled concepts (Fig. 30), I implemented the full system in two weeks, from detecting shape signals from the Arduino, collecting them in a queue buffer in Python, and detecting what people sang, collecting that in another buffer in Python, to linking them together in a simple ML model.

Once I turned my concept towards tangible and explainable AI, I started to reflect on the very first "AI" system I learned: a fully connected network with one hidden layer. It was simple enough to write out mathematically, and the performance was good enough to classify the gestures of the user.

The system to identify the pitch the user sang was from a straightforward tutorial for making a guitar tuner (Zurstraßen, 2020). The pitch was then converted to a note from A to G# dropping the relative octave information.

The labels were these musical notes a user sang, and the data to classify was the Arduino signal that captured the user's gestures on the sticky pouch. The network learned to assign the shape signal a

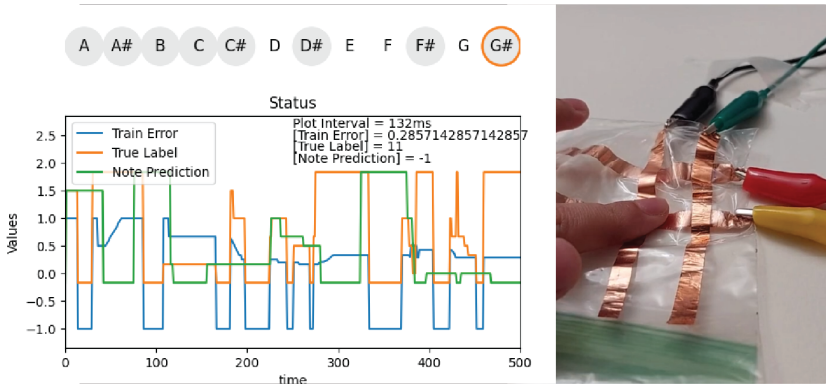


Figure 31: Left: visualization of the training process. The music notes are displayed above, and orange denotes the note I was singing, and green denotes the note the machine was predicting. Values are note numbers divided by 6. Right: The test setup of the pouch.

musical note that matched what the user sang in the training. The process is displayed in Fig. 31.

> Some tangents here: I was proud enough to use my own code from the course *Introduction to Machine Learning*, and ended up with many NaNs due to the instabilities in sigmoid activation functions that I naively implemented. Then I switched to PyTorch [5] reluctantly, thinking in dismay that this would make the system less graspable and explainable to the users.
>
> Anyway, I nearly abandoned the on-body wearable and the pouches after investing a lot of energy into designing an explainable AI system that "sings" with the user.

[5] PyTorch 1.13.1+cu116, https://pytorch.org/

One interaction episode looked like this:

1. The user touches the pouch sensor while singing a note (A-G#).
2. The ML model classifies the sensor signals as associated with that note.
3. User stops singing, ML reproduces the note.
4. User registers multiple notes to different touch gestures.
5. ML reproduces multiple notes when these gestures are repeated.

> An explainable AI system did not seem to have to be embodied. It was also hard to argue why an embodied and explainable AI system would result in music performance if the argument is to exploit people's own labor. Shouldn't the people feel the dread, the soreness, the anger of how hard ML was to train without corporate efforts? I even thought about making some kind of fitness equipment and gamify the experience into something you always lose.
>
> Let us go back to the concept of music performance, although it was already born with ominous signs. For ease of testing, I also use musical notes as the output of the model, so I would know if the model was predicting the same note that I sang.

The goal was to have the predicted output of the system match the note the user sang when they performed a touch pattern. One

indicator that the system was working correctly was that I could see three distinct outputs in the time sequence that corresponded to three gestures I provided.

> *[1/23] I added sound after I saw that the machine could produce 3 individual segments without my guidance, because I wanted to know if they were actually accurate.*

I sang extensively in my bedroom and successfully trained myself to sing well along with the gestures I was performing.

> *[1/23] My roommate told me she heard me singing all night.*

Eventually, the ML model was also trained, although it might be easy to train to begin with if not for my poor singing skills and ability to memorize what I sang. To make the output more fun and performative, which was among the earliest preferences of me to create sophisticated outputs, I made several attempts to make the machine prediction more interesting. I tried single sinewave, which sounded very thin and mechanical,

The first attempt to add music to the system was through the most basic `sounddevice (sd)` library [6] in Python, in combination with a sine wave array generated by the NumPy package[7], `np.sin(x)`. This rudimentary method did not achieve a satisfactory result, as a plain sine wave was missing the richness the overtones could offer.

[6] sounddevice 0.4.5, https://python-sounddevice.readthedocs.io/en/0.4.5

[7] NumPy 1.23.5, https://numpy.org/



> *[1/23] I scraped some code using the bare* `sounddevice (sd)` *library to output sinewave of a specific frequency with NumPy arrays, np.sin(x). Yes, the machine was following the notes I sang, often/half of the times one octave below.*
>
> *But the sound was awful. It sounded bare and thin and mechanical, (this will be a metaphor challenge) I do not have recordings of this, but it should be fairly easy to reproduce. It wasn't even pleasant to hear despite the joy I felt from my progress, because then I realized how much work I still needed to do to enrich this experience.*

and turned to Sonic Pias suggested by D. Byrne and Yanwen Dong.

In the next attempt, I used Sonic Pi[8] to add textures to the single note I was generating. It was fairly easy to pick up, because it was script-based in Ruby style. There were a wide variety of synthesizers to test with, and the effect improved greatly from a single sine wave.

[8] Sonic Pi 4.3.0 https://sonic-pi.net/

> *[1/23] [D. Byrne] and [Y. Dong] brought up Sonic Pi, and I spent a night reading the tutorial and got Python to send notes to Sonic Pi, when they were predicted.*

There I played with different synths, or different voices the machine could have,

> *[1/23] The piano sounded fine, but it was funny because Sonic Pi played individual notes each time I sent it, so I had to make the attack longer to blend them in. (Later I was able to use one note with long sustain to run, pause, and change it.)*

To make the performance more expressive, I consulted Zhenfang Chen, who suggested that I could compose the piece by choosing a theme and bringing associated musical elements or effects to the machine's prediction. Apart from special effects, which I did not have experience with, I chose to add background chord progressions that were labeled as "ethereal" based on a website.

[9] I had a tendency to write question marks in the journal when I did not bother to look up words.

and even added some background chord progressions scraped from the web. The chords I aimed for were supposed to have an "ethereal" effect,

> *[1/23] Then I thought it was still boring. [Z. Chen] suggested that I select a theme, collect some music elements based on the theme and play with them. So I randomly searched for "dreamy" chords, and got a page with "ethereal" chord progressions. [...]*
>
> *So I wrote the roots(?)[9] and qualities(?) of all the progressions, and made the machine find the pattern that has the chord that contains the currently predicted note, and play that chord, and the next ones as long as possible from the same progression pattern.*
>
> *[...]*
>
> *Okay, after all these, I made Sonic Pi align the incoming notes to the beats with a soft bass beat, play the notes on beat and discard the intermediate ones. Then I realized that because we are discarding so many notes, about 90% of them, the progressions did not make any difference at all (the sequence is mostly lost) so it was nearly random.*

which turned out to be a little bit like a haunted house effect. It successfully influenced my perception of the system. When I sang to my device, I could only sing gloomy and moody tunes then.

[1/23] Strangely, it still conveys a very strong vibe, despite the randomness and low volume, when the chord track was on, it was very hard for to deviate from 1) the note that it's playing, due to the tendency to harmonize? 2) the overall spooky, dark, weird mood, even when I want to sing some brighter melodies, like C E G major triads would be difficult. Maybe it's just me that get easily swayed into the background music?

Why would these chords be so magical?

I had a lot of trouble with recording the touch pouches, which were not inflated and not on-body, me singing, and the machine sounds. In what I achieved to show Professor Jesse Stiles, I was singing and tapping on a piece of dubiously looking, partly shiny piece of plastic fabric laid over my arm in a very small frame on the upper right corner of the video, and the main screen showed floating curves and bar charts as diagnostics (Fig. 32),

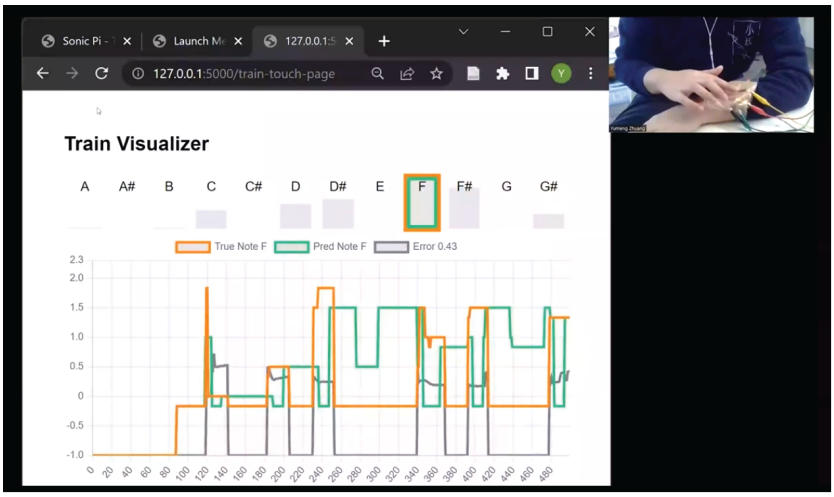


Figure 32: What I showed to Professor Jesse Stiles.

and a distant humming of synths and chords could be heard with great care. It took me about 15 minutes to explain what this mysterious Zoom recording was about, with all the concepts of hidden labor in ML, intimate musical performance, and so on. Eventually, he arrived at the conclusion that I was trying to make a digital musical instrument with an unconventional input interface.

In reality, when I showed the recording of me interacting with the classifier pouch to J. Stiles, he concluded that I was trying to make a musical instrument with pouches and machine learning. It did not occur to him that there was any effort involved. Instead, he interpreted the system as an attempt to bring a novel form of interaction to music-making.

This was of course very disappointing. Making a musical instrument did not touch on the speculative part, which according to me meant being something unusable. Nor did it touch

upon the concept of effort, because it only took me 3 seconds to register the shape signal to the note I was singing.

My observation through testing the system was that the ML classifier could learn the shape I provided within 3 seconds, and the difficulty was in the singing of the notes and remembering what I sang.

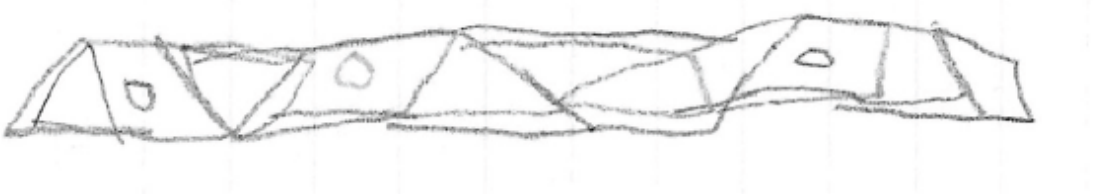
[10] Max MSP 8.5.2, https://cycling74.com/products/max

He did not grasp the necessity of ML anywhere, because the touch signals could possibly be analyzed with hand-crafted rules with ease. I had to agree with him that I needed to work on the argument a little more. He did introduce me to a helpful software Max MSP[10], which was more powerful than Sonic Pi in generating music, and I switched to Max MSP soon after this conversation.

Maybe because J. Stiles was a music professor, he did not perceive any effort in ML *or* singing, but when I presented this work to Malika Khurana and other students, they all perceived the difficulties in singing to the machine. To be honest, it was already hard for me to sing and hold the same note in order to register it with the shape signal, and "playing" this instrument in public was just impossible.

> *[1/23] For example, I could be singing C when open, E when touch one, G when touching two. Then when I want to reinforce C, I go back to open, I would not remember what note I associated it to. I would look at that flat segment of the orange line and guess, even with a later visualization of notes with name, I am lacking the ability to see "C" and sing it. So whenever I attempted to retrain the pattern to the same note, there would be an initial slide in my voice trying to slide the orange line to where it previously was. Sometimes this was enough to overwrite the dataset buffer the ML is training on, and confused the machine to think that this pattern (touch) belongs to another note. I would then be frustrated. Especially after I visualized how fast the training dataset gets filled (and gets depleted). I had to remember the pitch! Is this a major setback in terms of usability, or is it a "feature"?*

The system itself did not even work live anyway.



At a thesis mixer event, I showed people the system.

[1/23] which of course was not working, with me singing

People were self-conscious about their voices. People were not remembering what they sang and confused the ML algorithm. I should probably call this encounter of my thesis *a difficult music tutor*. Several people suggested that I try easier sounds,

M. Khurana said she would not be able to sing as well and as accurately as I could. She said what if some people were not good at singing, and if some people were self-conscious when singing. The singing challenge was really not something I wanted to challenge people with.

such as environmental sounds, and maybe speech, and, although I was not sure where this road would lead me,

Discussion: Hidden Labor in ML

I drifted from language generation and body gesture mapping in the previous encounter and discovered that controlling the shapes required a lot of ML training efforts. I interpreted the required efforts as an illustration of how machine learning models often exploit human labor without people understanding the process. I proposed that my artifact could provide an explicit experience of mental and physical labor involved in training machine learning models, elicit intimate feelings towards an on-body system unique to oneself, and foster conversation about the deep implications of this technology.

In summary, with the completion of the pouch-ML pipeline, I saw the opportunity to engage the user with a critical conversation about the hidden human labor in large-scale machine learning systems. From my personal interest in raising critical questions, I was not satisfied with the concept of a musical instrument with no effort while having most effort in singing, so I drifted away from ML that was ironically too powerful to "not train". I could imagine that introducing artificial complexity in the ML model or the pouch structure could bring back the effort, but that sounded too contrived.

I implemented the sound segmentation algorithm.

I was most inspired by the feedback M. Khurana gave me: It was too difficult to hold a music note and go back to it accurately, and she thought that recording longer clips of sound could be more engaging. D. Byrne provided me with similar feedback that environmental sound could be a rich source of input to a sound system. Therefore, I implemented the components in the following section.

The algorithm was simple and brute-forced and I seriously suspected that there were better standard algorithms,

Sound Implementation

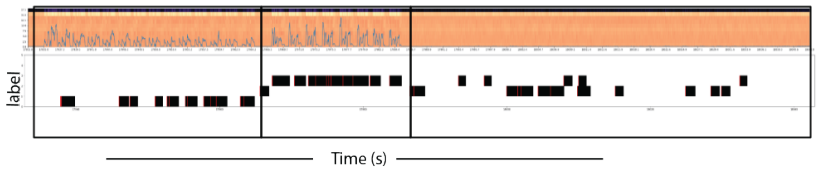


Figure 33: Audio segmentation. Box 1: identified sequence 1; Box 2: identified sequence 2; Box 3: sound interference.

Sound Segmentation To segment repeated sounds from a stream of input, I converted the audio signal to Mel-Frequency Cepstrum Coefficients (MFCCs) (Wong & Sridharan, 2001). On every 9-second sequence of MFCCs, I was able to use cosine similarity measure to locate repeating sequences between 1 second and 3 seconds long by comparing every pair of subsequences. I saved the unique sequences into a Python dictionary for retrieval. Fig. 33 shows one example audio clip.

but to fit into the story of graspable ML, I implemented my own. The testing and debugging forced me to repeat a lot of sounds: from random melodies to full phrases. All day I talked into the microphone on my earbuds and tested how the algorithm could identify and crop out the repeated phrases or melodies in the audio stream. My roommate said she heard me singing all night. This kind of labor was again unexpected. The moment I heard the mixed sound played back to me like a much delayed but everlasting echo, I was amazed by the richness of associations I had.

The first couple of tests were me humming some short melody, so the playback was simply like a chorus of multiple versions of me. Then I started testing with word phrases, which the system was also able to record and mix. Having phrases of varying loudness and lengths looping in a mix was a very surprising experience for me.

> *[2/7] [...] which makes me extremely self-conscious and embarrassed (?) by the sound I repeat. Therefore I prefer humming to words, but also short complaints and humorous remarks about my state of life.*

... such as, “Almost an idiot (in Mandarin)”, “Dreaming“, “What can I do?”, “Something else”. Hongyu Mao described it as “a cassette tape being stuck looping” and “acid style”, and J. Liu described it as an audio version of palimpsest.

Gradually, it became the most dominant feature of interacting with my pouches.

> Yes, I still had the pouches, because I convinced myself that I would like the pouch to stay in specific shapes without relaxing to the boring resting state, which required sticky patches. As for what they would look like and where to put them, I was not sure.

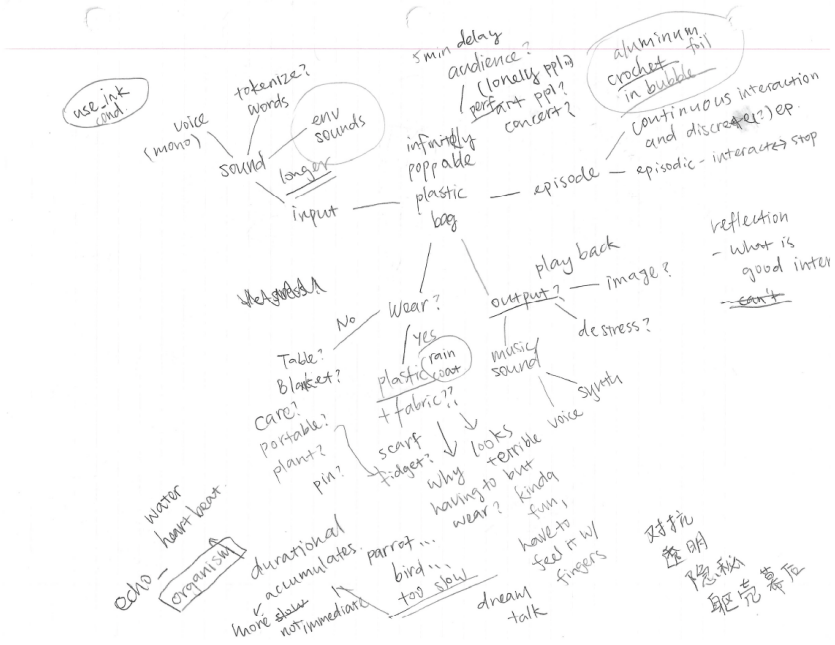


Figure 34: Concept mapping stemming from the poppable plastic bag.

Moving from the concept of human labor in machine learning, I developed the following goals for my interaction:

- Comprehensible, understandable, not perceived as random but has a cause-effect logic
- Slow and full of anticipation as required by the air pouches
- Flexible and programmable, supports improvisation
- Takes care and effort to maintain
- Low bar/skill to interact with

I include a sketch of concepts stemming from the poppable pouches in Fig. 34.

> The domination of repetition was intriguing to me. Sound segmentation also made the idea of ambiguity alive again, because I really could not say any single note was ambiguous or mix them in ways I could call ambiguous, but I could now mix the phrases and they would sound fine when I play them back aloud. The mixing of phrases was realized by weighing the gain of the recorded clips with the outputs of the last layer of the ML model. The ML model as a whole still learned the mapping from pouch shapes to sound phrases but was now giving a list of confidence scores across the clips.

Machine Learning Implementation

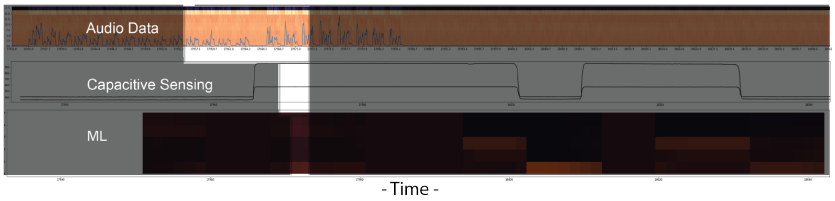


Figure 35: Illustration of how audio data and capacitive sensing feed into the ML model.

In the classification and retrieval stage, I linked the incoming queue of Arduino signals with the sound segments through an ML classification mode (Fig. 35). At training time, the user's phrase was recorded along with the shape signals. Then, for each segment of the input shape signal, the model predicted the index of the corresponding sound segment stored in the dictionary and the relative weights of each sequence. The shape signal buffer was constantly filled with new data, as well as the sound library.

The predicted weights of all sequences were used to mix the sound segments in the dictionary to achieve a chaotic effect. The mixing was produced in Max MSP via a User Datagram Protocol (UDP) protocol that passes the weights from Python to Max MSP.

> Would a person say a mixture of clear phrases become ambiguous? This is highly debatable, but I let it be. What I did that was not as ideal here was that I equated mixing with chaos, which is something I still do now, but worse, chaos with forgetting.
>
> The concept of forgetting came from the effect of mixing phrases and also the artificial forgetting I included in the backbone ML model which associated shapes with sounds. The new shapes filled the training buffer constantly, and the model would "forget" earlier shapes, making way for the newer shapes.

There is also an important concept I introduced called forgetting. At the previous encounter, while the user provided one of the 12 musical notes and a touch gesture, this data pair was stored as training data in each of the note "buckets" with a fixed size. However, if the user did not go back to the note for a while, the training data in the bucket would be stored for a long time, and the user would not be able to remember or reuse the touch gesture stored in the bucket. Therefore, in order for the gesture to be reusable, I artificially dropped training samples at each call of the training process, so the system unlearned unused notes and gestures.

> The library of sound clips was also limited to 5 phrases with new clips overwriting the old ones.

For this encounter, the 12 musical notes were replaced by the 5 sound clips. The constant replacement of sound clips led me to the concept of digital memory. The words I said would momentarily be remembered, be repeated, and be forgotten, well before I forgot them.



Now I think I was conflating two kinds of forgetting here: the actual deletion of the older shapes and clips of sounds, and the effect of forgetting obtained from the mixing of phrases. Whichever I called forgetting shifted and morphed as I wrote the arguments, but towards the end, I landed on machine forgetting, which was both of these, and human remembering, as resulted from the vocal repetition of phrases.

Remembering through Repetition In psychology, it is a well-known phenomenon that we remember better, and for example, recall more details of autobiographical memory when the memory is retrieved more times (Nadel et al., 2007). The device from this encounter asks for many repetitions of the sound or phrase the wearer wants to record, thereby paradoxically enhancing the wearer's memory of the feeling or event while being a forgetful device.

Forgetting in Digital Age It is hard for machines to forget. Mayer-Schönberger (2011) observes that it has been easier than ever to record information in digital forms, due to the dramatically reduced price of storage and improved interface design. Forgetting, on the other hand, is not built into the physics of digital storage. This encounter utilizes the uncertainties and instabilities in machine learning to create a natural process for machines to forget. Furthermore, the shape-changing interface that the machine is learning is naturally susceptible to shifts in configuration, erasing its previous states.

In the scope the device I propose could remember, it is only a companion witness of the life of the wearer, not a complete capture of their experience. Sellen and Whittaker (2010) emphasize this point that machines could only provide cues that "trigger different kinds of memories". My work is enriched by the concept of cues as presented in the Memory Device by Bertran (2014), which only records timestamps of events the user chooses to remember, and displays them as simple lines (Fig. 36). My device stores these records in its deformation which is not explicitly legible, and the audio recordings can only

be cues as well, being limited to 1 – 3 seconds.



Figure 36: The Memory Device, credit Bertran (2014)

Landing: Fallible Digital Memory

We tend to record things that are slightly out of the ordinary. Digital storage is so affordable that it is essentially unlimited for many people, and the effort in recording data into digital storage is also negligible. Digital media as a form of memory storage are precise, accurate, and immutable as well. To some extent, this effortless memorization encourages an illusion that life is full of extraordinary events we can relive with ease. Through an interactive wearable, this thesis challenges this illusion by making the events less extraordinary at the time of recording and by making memory storage scarce and unstable, thereby enhancing the individual's connection with the event.

More specifically, my thesis raises the following questions:

- What kind of friction and benefit would repetition create in an interaction with a memory storage device?
- What adaptation would scarce and unstable memory prompt an individual to adopt when achieving meaningful interactions?
- What feelings would a large piece of wearable evoke as it isolates the individual from the external environment?

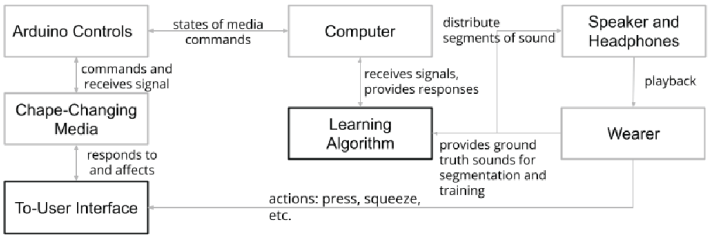
My thesis provides insights into these questions through the development and evaluation of the Transparent Puffer Coat, a wearable system that uses real-time signal processing and machine learning (ML) to achieve sound and shape-based information storage and retrieval (Fig. 37). The Coat serves as a voice recorder and the user has to patiently repeat each phrase many times while wearing it to record them. It stores the phrases in the shape deformation of the Coat and plays back with uncertainty a mixture of what it has recorded based on its current deformation. It has limited storage of phrases and new phrases overwrite the old ones.

The four components of the Coat, the ML model, the shape sensors, the sound processor, and the wearable structure, act together to spark reflections and conversations around the influence of technology on our habit of noticing and remembering.

The main contribution of this thesis lies in the following areas:

- An interactive system for sound recording and retrieving:
  - That senses touch gestures through inflatable pouches.
  - That communicates information based on shape changes and audio feedback.

Figure 37: Components of the Transparent Puffer Coat.



- An extendable pipeline consisting of online training of a neural network for classifying shape deformations, the identification of repeated sound patterns in a real-time audio stream, and the mapping and update mechanism of shape-audio pairs.
- An on-body, touch, and voice-based intimate experience of an unconventional memory storage device.
- Proposed opportunities in speculative design for memory.

> I remember I had an urgent deadline for a presentation that required a video, so I briefly abandoned the pouches

Another line of inquiry was along the form of the air pouches that would fit the theme of memory. Some explorations are included in Fig. 38. I imagined that the system could be an AI system that requires love and attention like a small creature. I experimented with various forms of placements and materials.

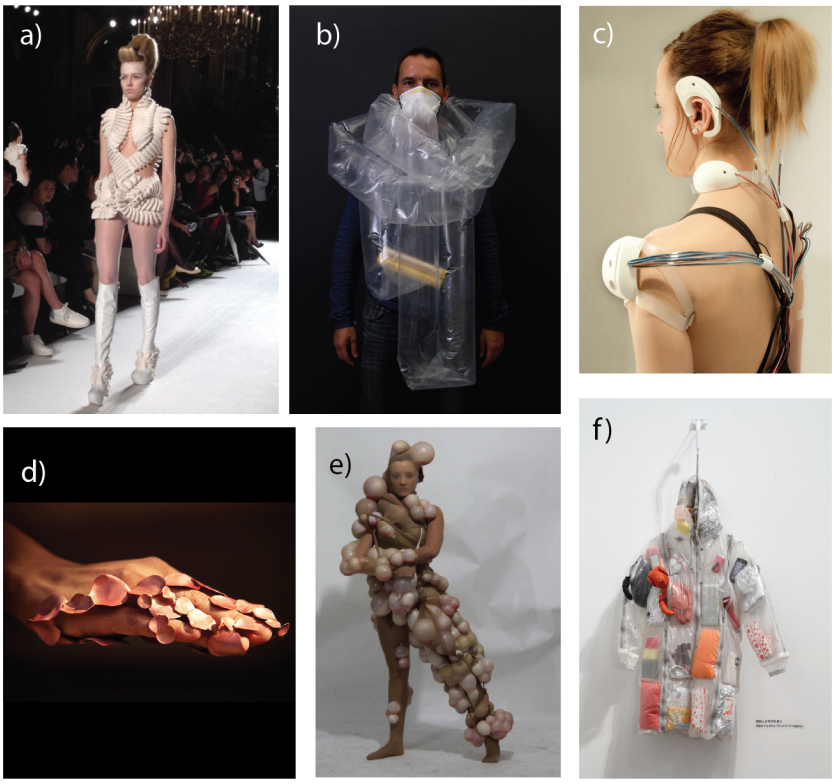


Figure 38: A few wearable precedents that inspired the development of my wearable. a) Sculpted Cashmere, Fang (2013), photo credit: Quinn Connors b) Air Slaves, Let it Be! art agency (2017) c) mediated reality, ling QL (2013) d) Intangible Embrace, Eriksson (2011) e) Freakshow, Birch (2013) f) FINAL HOME, Tsumura (2015)



Figure 39: Various forms that were inspired by the precedents.

[1/23] Not much wearable yet.

I tried taping the pouch around my arm, the capacitive sensing values were a lot bigger, but thankfully had not saturated.

It felt okay pressing into myself, but I realized that the machine was sensitive enough to pick up every "muscle" move, so I either had to train it while in multiple gestures or I had to freeze absolutely to get it to recall this note. Is this a feature, too?

When I put it under my sweater, this is even more obvious. My hands could only switch from pressing into the pouch (so that it bulged and pressed into the sensors) to loosely holding the pouch. I couldn't hold, let go, and hold it back to the original position...

Also the airbag under a sweater made it super warm. The wiring coming out of the holes in the middle of the sweater were an interesting sight (Fig. 41).



Figure 41: The experiment with a pouch under my sweater.

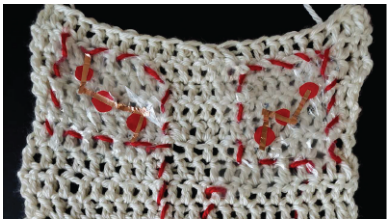


Figure 40: The experiment to sew pouches onto a crochet piece.

Should I expose or hide the pouches?

[1/29] I tried a crochet piece with two plastic pieces sewed on (Fig. 40). It looked like some kind of patchwork for holes/tears. The red yarn does stand out, and perhaps with careful arrangements would look fine, but I found myself overly drawn to the concept of craft, woman, coziness, etc. The pouches would be floating pieces of trash in a pile of wood. I made a pouch "scarf" as well, though only our pf plastic packing pouches, but still looked like a piece of trash.

and made an actual wearable with capacitive sensors and the working sound segmentation algorithm. I reasoned that I could simply raise my arm and hold it there as one shape, drop my arm to make the other, and record two phrases respectively.

The following section showed my tested implementation of the transparent raincoat. There is also a video showing the coat in action.

Raincoat Strips

> *[2/13] I had sensor/tape strips running parallel to my arm down the armpit and horizontal strips across my chest, to under the arm as well. All four strips were extended to the center (?) where the buttons line up (Fig. 42).*
>
> *It was extremely hot after 3 or 4 recordings and by the end I took off the coat, there was already condensations in the sleeves and my shirt was damp.*

I tested the coat and recorded a demo video, with two poses and two phrases: "Dreaming" and "Children are playing". The screenshot from the video shows the Coat in action (Fig. 43).

> *[2/13] I recorded a demo with the raincoat I got many days ago. [...]*
>
> *I recorded 6 times, because I did not realize that the resting state, with my right arm down, would take a while to register back (settle down) from the raised state, while the raised state could be triggered almost immediately, not sure why.*



Figure 43: Singing the phrase with the arm raised.



Figure 42: How the sensor strips are laid out on the plastic raincoat.

The wearable was made out of a transparent raincoat, which was inspired by a Japanese artist Kosuke Tsumura's FINAL HOME (Tsumura, 2015), and I loved it so much that I wrote an entire fictional "use scenario" of a lonely girl living with the raincoat which briefly remembered and then forgot her day. I was thinking about muscle memory, the body as a container, and the indelible scars and marks that store our life experiences implicitly, like the shape signals from the sticky patches. I even illustrated the story with a comic page, but nobody had the attention span to read the story in order to understand the comic (Fig. 45).



Figure 44: Illustration for the fictional "use-scenario".

Inspired by a transparent raincoat with pocket stuffed with various objects in Kosuke Tsumura's Final Home series (Tsumura, 2015), I thought the plastic pouches could be artfully integrated into the plastic features of a raincoat as well. Around early February, I wrote a story about my project as a transparent raincoat called TR who lives with a girl living alone in her apartment (illustrated with Fig. 44). The girl uses the Transparent Raincoat (TR) as an interactive tool for recording and remembering phrases. TR is connected to a machine in the girl's apartment and can repeat the phrases she sings into it. The girl activates different air-filled chambers in the TR by pressing on specific red dots on the fabric, causing it to recognize a new phrase. Throughout the day, the girl will enter the TR and sing new phrases or recall previously recorded ones by pressing the corresponding dots. In my writing, the girl treats TR with a sense of care and affection. Her gentle handling of TR, her routine of turning him towards her every morning, and her repeated singing of phrases suggest that she values TR and finds joy in their interactions. The full story could be found in Appendix B: A Day of a Transparent Raincoat by the Window.

> Anyway, the video was successful, and many viewers said they finally understood what I had been trying to do in the past months. At that moment, however, I was already wondering why I would want something digital that forgets, outside of the story in which it was merely a passive prop and an imaginary friend to the girl to render delicate loneliness and caring of inanimate objects. Somehow this felt very intimate again.



Figure 45: The comic that accompanies the raincoat story.

Discussion: Fallible Digital Memory

As machine learning turned out to be too easy and music performance was not for everyone, I drifted to manipulating longer sequences of sound and arrived at a mesmerizing effect that appeared chaotic but fully consisted of the user's memory. I argued that the forgetfulness and uncertainty induced by the machine learning model is an alternative narrative to the faithfulness we expect from standard data storage devices and would prompt people to reflect on their recording practice and their relationship to events they often carelessly throw into storage.

The first development for this encounter was the sound segmentation algorithm that processes human sounds. Then I linked the segmented sound to capacitive sensing from the pouches via a machine-learning classification model. Attempts were made to test the form of the pouches and how to put them on the body, from under the clothing to exposing the sensors. I also wrote a imaginary scenario that depicted how the pouches, attached to a raincoat, could function in a person's life.

> As I was preparing a submission to Design Interactive System Demo track, I read many papers from previous years trying to gauge what kind of argument they made. I became worried when I found that their arguments seemed to be very positive, contributing to society in tangible ways. Would a lonely person actually want a companion that forgets their day? That was more like a lyrical fantasy to me now.

When I was refining the argument of digital forgetting for a submission to the Demo track of Design Interactive Systems, I started to be bothered by two "flaws" in this argument. First, what was the exact effect of repetition? Did I remember the phrase better, and thus the memory of the event, or did I make the event less remarkable as I argued in the imaginary introduction, and forget with the unstable device? The physical effort of repeating seems to suggest the first, while the repetition and deletion seem to suggest the latter.

Another problem was that I had trouble arguing "so what". Why do we want this alternative narrative? Given the ambiguity of what it did to me, it was very difficult to argue for a larger social implication.

Nevertheless, this encounter brought me closer to a playful interaction of recording longer phrases. I have also formulated a theme around care and memory, two internal processes users would explore interactively with my piece.

> Upon drawing one of the many mind maps around the concept of human, remembering, machines, body, etc., I was drawn to the concept of "here and now" (Fig. 46 and Fig. 47). I wanted to emphasize the human-like qualities of machines that could forget, and how that

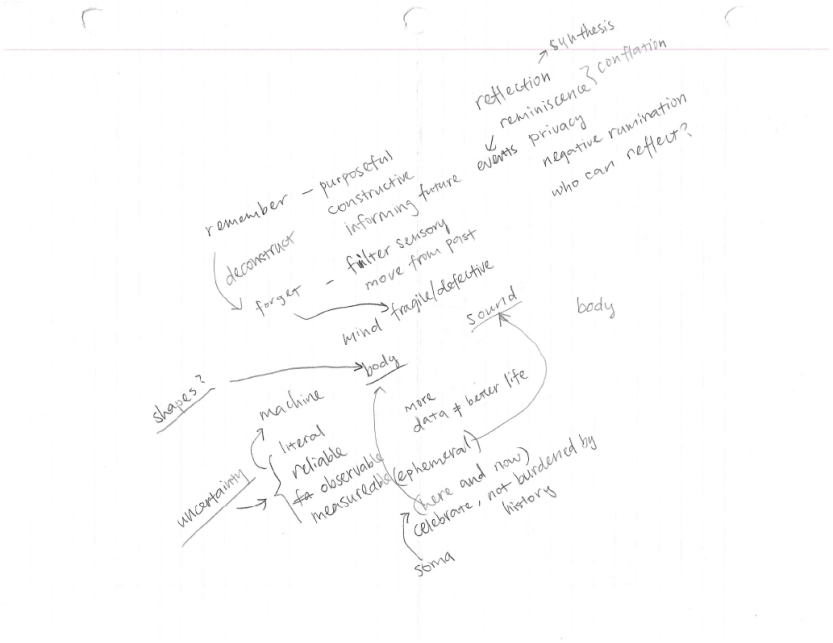


Figure 46: The concepts around body and sound which led to reflection and rumination.



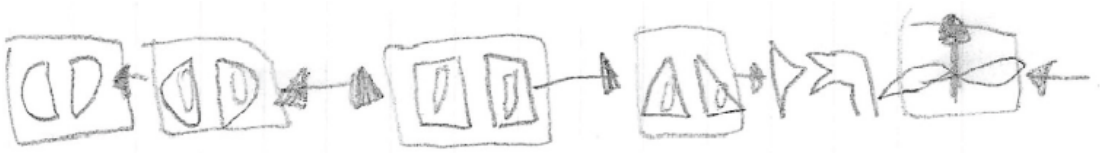
Figure 47: An enriched and refined version of concepts leading to the "here-and-now".

would help humans to live, and I might also be listening to some songs that followed the line of living in the moment. It was reflected in many articles dealing with machine forgetting as well. This was not very far from the argument of machine forgetting, though I did not find it suitable to call the mixing "forgetting" anymore.

Gradually pieces started to fall into place. The vocal repetition of phrases actually brought about another quality,

To address the confusion of what the repetition did to me as well as why we need a machine that forgets, and if it is actually forgetting after all, I tuned into the embarrassment, the awkwardness, the confusing, and haunting quality of the interaction during the development of the sound segmentation system. When I was singing phrases to melodies, I typically lost the focus on the phrase I was saying after 5 repetitions. My ears were also ringing after the long tests, with the waves of the sound.

already hidden in the written introduction of the previous proposition. I said that human remembered better as a result of repetition, but I also wrote that repetition made the phrase less remarkable at the time of recording, so the machine could forget and the user would not need it to remember anyway. This peculiar quality of repetition turned out to have a better explanation, which reminded me of mental health therapy where they asked people to repeat a negative thought towards themselves in silly ways and end up feeling less fixated on it.



These experiences reminded me of a technique I learned from mental health therapy: silly repetitions could disengage the speaker from the negative thoughts they have, either towards themselves or their situation. This technique was part of the numerous techniques under the term Cognitive Defusion.

Cognitive Defusion There is a vast amount of literature on Cognitive Behavioral Therapy and numerous techniques developed to address maladaptive cognitive processes (O'Donohue, 2009). Cognitive Fusion, as discussed by O'Donohue (2009), is a normal process in which individuals relate other thoughts to the event they are experiencing. They provide the example that a person fixing a car could mentally review the operation steps before starting. This process becomes problematic when a person's thought takes away other probabilities of functioning in the world, such as one's anxiety about public space

may cause the person to picture the public space as only stressful and neglect other potential outcomes. Cognitive Defusion techniques are techniques that detach the individual from their thoughts and inner events so that they can function with more flexibility.

One classic technique from Cognitive Defusion is asking the individual to repeat a phrase many times. Experience shows that by rapidly repeating a phrase, the auditory quality of the phrase itself will become more dominant, and the literal meaning of the phrase will be reduced, thus its influence on the speaker.

> I have to admit that I had never been a fan of this method, nor had I liked the idea of repeating things out loud at home in the first place, but the extensive testing of the segmentation algorithm really indicated that this would fit.

I disliked this idea very much the first time I heard about it, precisely because it felt very awkward, but the intensive testing prompted me to think of this as a design space.

> Again, I started to struggle with framing the argument: are the phrases being repeated bad or negative? Or can they be neutral?

Landing: Tangible Rumination

When we think of something repeatedly, we might be “ruminating”. Merriam-Webster’s Dictionary defines "to ruminate" as “to go over in mind repeatedly and often casually or slowly” (Merriam-Webster, n.d.). This seemingly counter-productive activity also “suggests little of either purposive thinking or rapt absorption”. Rumination in our brains contrasts with software and devices which can recall and present to us our recorded thoughts in a linear fashion as efficiently as a click of a button. However, while productivity is valuable, *purposelessness* has its unique role in life as well, such as bringing back playfulness and breaking the momentum of analyzing life. This thesis examines the potential for digital devices to mediate rumination instead of eliminating it.

Specifically, this thesis aims to bring about purposelessness through an interactive device that encourages users to repeat a thought and reduces the meaning of it to a playful level. Through this device, I investigate the following questions:

- How could the repetition of thoughts change their meaning and the speaker’s attitude towards it, as supported by a recorded technique from Cognitive Defusion (O’Donohue, 2009)?
- How would feeling one’s voice via a tangible device enhance people’s experience of the here and now?

- What factors should be considered in designing for rumination?

A physical prototype was made to facilitate the investigation. The Echo Mask is an interactive device consisting of a partially inflated mask and a wooden box. Real-time sound segmentation and machine learning (ML) algorithms were used to register the user's thoughts with the shape of the mask and to enable retrieving the recordings based on future manipulation of the mask. I utilize the uncertainties in ML models as well as the inflated pouches to create a chaotic, looping playback of the recordings while the user speaks into the mask. The inflated mask also allows the user to feel the vibrations their own voice produces.

Through testing and self-use, I observed that the meaning of the thoughts I spoke into the mask was greatly reduced to waves of background sounds. The feeling of the vibrations also made me aware of the action of speaking. Further, I identify three design considerations for designing for rumination: placement of the device, the user's self-consciousness, and the device's agency.

The main contributions of this thesis lie in the following areas:

- A physical prototype of the Echo Mask that explores technology-mediated rumination.
- An Arduino-sound-ML pipeline for leveraging physical and algorithmic uncertainties in computing.
- Three detailed design considerations that open up the design space of rumination and human-technology relationships.

> Tomas Cabezon Pedroso told me when I showed him the manuscript that nobody would tell a machine good news over and over again in order to feel less attached to them. He was probably right. Especially when I called the repetition rumination, which had a mental health connotation of negative looping thoughts,

Rumination Definition From the Merriam-Webster Dictionary, rumination has the following definitions:

> Ruminate as a transitive verb:
>
> 1: to go over in the mind repeatedly and often casually or slowly
>
> 2: to chew repeatedly for an extended period
>
> intransitive verb:
>
> 1: to chew again what has been chewed slightly and swallowed: chew the cud
>
> 2: to engage in contemplation: reflect

In Merriam-Webster Medical Dictionary, rumination is defined as: the act or process of ruminating:

> a: the act or process of regurgitating and chewing again previously swallowed food
>
> b: obsessive thinking about an idea, situation, or choice especially when it interferes with normal mental functioning. Specifically: a focusing of one's attention on negative or distressing thoughts or feelings that when excessive or prolonged may lead to or exacerbate an episode of depression

it seemed a bit difficult to argue that the benign version of rumination, or mindless repetition of thoughts, could be a design opportunity for digital devices.

Reflection and Rumination in HCI Researchers of systems that "capture" our life critique the approach that only focuses on technical breakthroughs instead of carefully evaluating and stating the benefit of their innovations (Sellen & Whittaker, 2010). To bridge the gap, Sellen and Whittaker (2010) suggests that devices that store life data could provide the "5Rs" as benefits": "recollecting", or to remember past details, "reminiscing", or to re-experience past emotions or sentiments, "retrieving" of information, "reflecting", or to examine the past experience to learn about self, and "remembering intentions" of activities and decisions.



Zooming into "reflecting" now, Baumer et al. (2014) define reflection as "reviewing a series of previous experiences, events, stories, etc., and putting them together in such a way as to come to a better understanding or to gain some sort of insight". The authors located several areas where reflection is considered positive, such as supporting behavior change, increasing self-awareness, or encouraging better educational outcomes.

There are works that specifically target the differentiation of reflection and rumination. Despite the similar positive effects observed above, Mols et al. (2016) note that some people simply want to reflect less. Treating rumination as a negative version of reflection,

the researchers call for caution in designing for everyday reflection and attention to "constructive reflection". Similar caution is reflected explicitly by Eikey et al. (2021), as they define rumination from the field of psychology as "anxious attention paid to the self, often tinged with a fear of failure and having detrimental impacts on self-worth, which negatively affects self-improvement efforts". This argument is based on the earlier psychology research by Trapnell and Campbell (1999), in which the authors consider both reflection and rumination as under the concept of "self-attentiveness" while rumination has a negative focus, closely aligned with the definition in Merriam Webster Medical.

> Maybe I should go back to the physical artifacts, as these arguments were all in my head. After recording the raincoat video, I struggled with making a wearable piece with pouches, because the weight of the coat itself would pull the pouches back to their straight and flat shapes regardless of the sticky patches.

I did not, however, put air pouches onto the raincoat, because the raincoat was too large for the air pouches to not deform when the user moves. For example, if the button pouch is pressed down in the center, any stretch sideways would most likely pop the button. The weight of the raincoat would be pulling the pouches down whenever the user stands up, making sound retrieval difficult to control.

> *[2/18] I want to make a pouch that changes shape more dramatically than a single pressing like a pillow, because a coat is a large piece and simple buttons will be hard to see/less expressive. Also, another issue is that a coat deforms greatly when taken off, so if I want to fit it into the narrative of a lonely repeating speaker I have to somehow maintain the pouch's contact points.*

Pouch Implementation

I brainstormed a few pouch structures I had encountered in L. Yao's and *Kinetic Fabrics* classes, taught by Professor Olivia Robinson and Dr. Garth Zeglin. In that particular moment, I considered making something with obvious movements, to compensate for the scale smaller than a raincoat. I knew that pouches contract when they are inflated, and pinching into the pouch from one end was equivalent to inflating the other side. The most successful attept was the curling tail, which was a triangular pouch with parallel partially sealed lines spaced out across its body. The tail curled up when the larger end of the pouch was pinched and the tape stuck to the other side.

> *So I basically made a contracting curling tail, a long triangle flap with reinforcement on one side so it curls towards the other side when inflated. The sticky flap is*

put at the larger end of the long triangle such that when pressed down, the air will be pushed towards the thin body with partially sealed lines, as shown in Fig. 48.

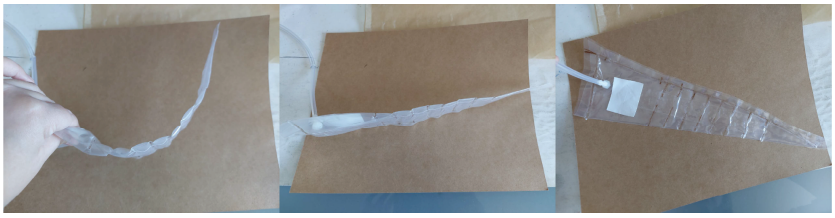


Figure 48: The curling tail pouch. Left: the pouch was pinched. Middle: the pouch was relaxed. Right: the seal pattern of the curling tail.

In terms of where to put the pouches... I find that putting it on the body and pressing directly into the body is actually not enough contact force to secure the tape contact, especially on places where I have a lot of fat, not even forehead works (Fig. 49). I have to pinch with my thumb and index finger.

I guess a hat is easier to make but a coat is more interesting.

It was also very difficult to make large inflatables in general, even after a much-belated realization that many household materials were heat-sealable. Combined with the argument of rumination, I thought I could make an inflated hat because thoughts come from my head and the act of repeating them is a metaphor of releasing them into this alien, almost parasitic, device (Fig. 50).

Relating the concept of rumination to a person's thoughts in their head, I started to design other different head wear, based on the curling mechanism I discovered. The artifact at this stage had a short name: the Cap.

I drifted away from the tails because the force they required to hold the shape was so strong that it was hard to get them into the desired shape by pressing into the body. Out of the two other options in Fig. 50, I landed on the "helmet" cap, because it could provide a degree of isolation, as if the wearer is left alone with their thoughts. I manufactured the helmet cap out of large thermoplastic sheets. The cut patterns came from a freely assembled hat with irregular paper pieces I made. The next step was to put the tube fittings through one layer of the pouches. I then used a soldering iron to seal the edges of the pouches. Through out the process, I marked the pouches and tubes with labels to keep track of them.

[3/4] Last weekend (2/26) I tried to make a hat. I cut random paper pieces and put them (taped) over my gallery hat (the hat I got from National Gallery of Art), which were 14 of them. Then I realized that I don't have that



Figure 49: Taping one curling tail to my forehead and overlaying and mirroring many photos.

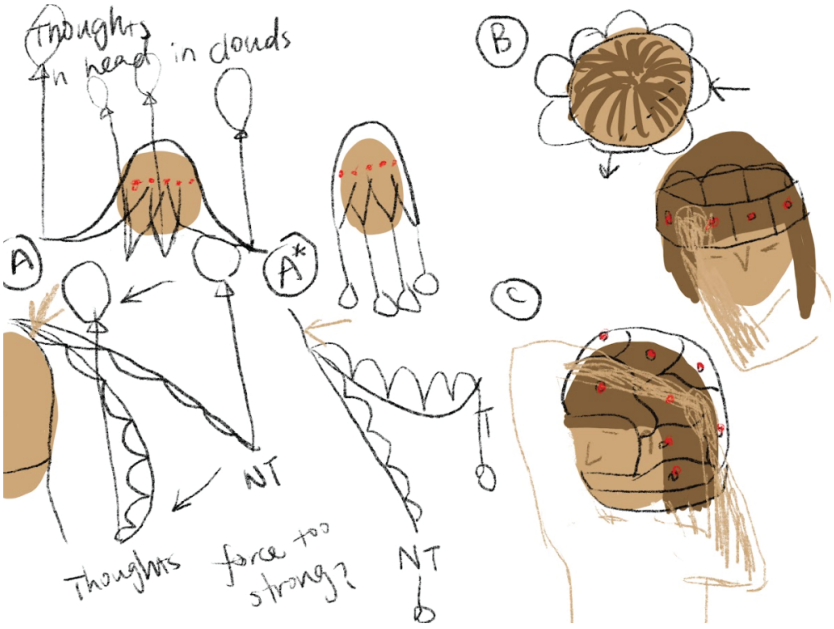


Figure 50: Some early thoughts on head-mount wearables.

many connectors, T connectors, probably have only 8. So I selected 8 of them to be inflatable. I cut the labeled paper pieces off and used them as patterns to make the pouches (Fig. 51).

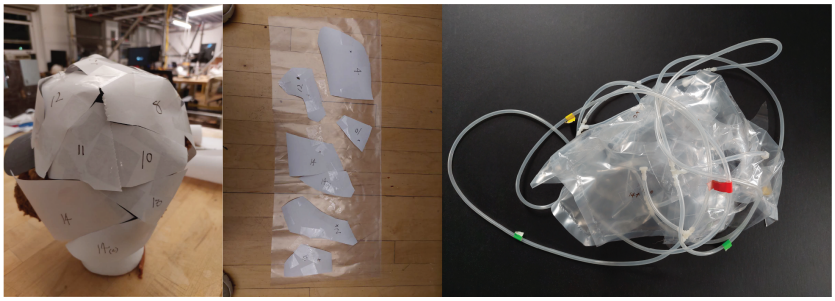


Figure 51: The (disappointing) process of making an inflatable hat.

I did make a hat, which looked like a pile of trash,



Figure 52: What the "hat" looked like on someone's head.

Unfortunately, I was not pleased with the result of the helmet cap. The pieces that aligned perfectly in Fig. 51 contracted and distorted when I inflated them, despite the margins I left. In some places, the margins were too large and I could barely recognize the pouch for where it should be. It was nearly impossible to assemble the pieces back to a spherical shape.

[3/4] It was not successful: the pouches contract when inflated, so they don't conform to the curved surface.

Also, I added too much margin to the pouches such that they don't fit together like a 3D puzzle any more.

Overall, it looks like a random pile of white plastic with semi-transparent tubes. Kind of like a failed chemical experiment set. It is very hard to imagine putting it over my head. When I did put it over my head / as well as a foam model's head, it looked a little like a veil (per Howard Han) with brain-computer interface (Fig. 52). The packing tape is floating everywhere.

and D. Byrne's question of why I wanted to make the interaction out of sight ended up killing this idea. I asked ChatGPT about different forms the rumination device could take, given the vague introduction similar to what I included earlier.

After I wrote a few paragraphs similar to the Introduction above on tangible rumination, I fed the writings into ChatGPT[11] and asked for ideas on forms other than the "Cap". Overall, it was able to comprehend most of my writing, and reasoned that the Cap would be some kind of novel interactive device. Following its understanding, I purposely asked for ideas on other non-wearable forms the Cap could take.

[11] OpenAI's ChatGPT (https://openai.com/blog/chatgpt)

[3/4] I asked GPT why and why not the Cap should be a wearable. Some points are good: Wearable: Hat, headband

- *Natural, embodied experience*
- *Movements and gestures*
- *(debatable: intuitive) and immersive*

Non-wearable: furniture, fixture, vase, bowl, shelf/case (arrange items?? Interesting)

- *Discreet, inconspicuous, subtle, non-intrusive*

The "fixture" category inspired me, especially the bowl idea. I imagined it used when a person spitting thoughts into a bow and looking at them, or dumping them out.

ChatGPT seemed to be supportive and creative, and I filtered out several inspirations: if I allowed the experience to become not embodied, I could work on the form of it to match its purpose, and fixtures such as vases and bowls would be good options. The container shapes are friendly to make, and they could symbolize the temporary place thoughts could stay.

This led to the current form of the Echo Mask (Fig. 53).

[3/4] So I taped 4 pouches together and clipped the mic in the center hole. I buried my face into the bowl, half inflated, and sang into the microphone.

The pouches fogged up a bit as I sang. I could feel the vibrations of my voice with the fingers pressed



Figure 53: The Echo Mask in its current form.

outside the inflated pouch. This feeling is quite interesting. I could still see the outside as I sang. But it wasn't as clear since there's two layers of plastic. I like this semi-private/secluded feeling. I plan to make a base for this device to rest on, and have a speaker in the base, so the "private" thoughts will be broadcasted.

I also made a box out of 3mm plywood sheets of the size of a small Amazon packing box. It had several rectangular holes on the top for the speaker sound to go through, a circular hole for tubing, microphone, and wirings to go through. The Arduino and the speaker were placed inside the box. It turned out that having the speaker right next to the microphone was problematic for the sound segmentation algorithm to work, but I did not implement a solution yet.

Now I'm about one month into submitting the manuscript of the Echo Mask,

Discussion: Tangible Rumination

In this encounter, I propose that digital devices could mediate rumination in a tangible way by encouraging users to detach meanings from thoughts and letting them feel the vibration of their own voice. The Echo Mask, the artifact resulting from this encounter, consists of a mask made out of partially inflated air pouches with sensors and a laser-cut wooden box holding the hardware components. A machine learning and sound segmentation pipeline registers the shape of the mask and links it with sound recordings that the user repeats into the mask. Through repetition, the user would experience reduced meaning of the phrase, as well as a mixed, chaotic playback of all recordings.

I already start to worry about some of the inconsistencies in the argument as well as the artifact.



However, I found it difficult for me to actually use the technique of

cognitive defusion in its original form, that is, to reduce attachment to negative thoughts. I wrote:

> *[2/26] Since I've tapped into the idea of rumination, I've been thinking about things (negative) about myself I can use to test the system its ability help me "defuse" form the negative thoughts. However, the mere thought of speaking uncomfortable / hurtful words aloud, even to a machine in my room, makes me anxious and uneasy. I think I'm going to back track and reduce the idea of cognitive defusion.*

The backtrack did not fully happen, but I still tweaked it in the neutral direction. I think this difficulty hints at the machine's agency over the user, as it is listening to the private thoughts, distorting them and playing them back to the public, without full controls from the user.

> For example, was machine learning really needed? The mixed playback could be achieved by some hard-coded logic, and why would people need to store the phrase they are repeating or ruminating about and retrieve them with the associated shape? Why do I need inflated pouches to hold the shape, when any rigid material like buttons could also do it?

This is the last encounter of the development of the air-pouch-based interactive system. It is very tempting to move forward from here still. Am I happy with the tons of analogies and metaphors involved, from repetition to rumination? Am I happy with the placement of the microphone inside the mask, which interferes greatly with the recording? Is the mask intuitive enough to interact with, asking users to put their faces in? I will probably say no, not sure, to many questions.

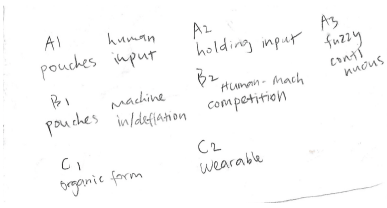


Figure 54: Some core elements I wanted to keep that I developed while working on Hidden Labor in ML.

Let me go back to these elements in Fig. 54 that I had wanted to include at the encounter with *Hidden Labor in ML*: A's are about the pouches, B's are about the machines, C's are about the form of the device. It is very tempting to find a universal theory to unite the elements. It has not happened yet.

> I guess I am simply too greedy, trying to find a universal theory to bundle all my past attempts into. In some sense, this bigger thesis as it is now is my solution to the craving for unification.

Conclusion of the Journey

Krogh and Koskinen (2020) discussed four sources of knowledge: artifact, application of the correct, often rational and scientific, method, discourse around the artifact, and insights from user participation. In my development of the four encounters, I sought to generate knowledge from these four sources.

The first source of knowledge is the pouch system itself. As I introduced in the Background chapter, Carroll and Kellogg (1989) describe that design artifacts hold knowledge as they are "theory-nexuses". In my work, I always had air pouches as the center of the design. The air pouches had inherent qualities, as I observed in the early encounters, that the signals from them were unstable. However, the unstable signal is not knowledge. It was the assembly of the system, for example, the mask and the box, that gives the artifact a story to tell. This knowledge would not be possible without implementing the system.

The knowledge generated from the discourse of the artifact was perhaps the most important contribution I was trying to make in the development of the artifact in the four encounters. One design artifact could support multiple readings, a property the Krogh and Koskinen (2020) termed as interpretive possibilities. The collection of readings, if it contributes to a greater dialogue in knowledge production, is knowledge as well. In *Intimate Ambiguity*, I explored how the ambiguity in sensing, body gestures, and self-perception could be expressed in ambiguous language. In *Hidden Labor in ML*, I questioned if feeling the labor of training an AI system unique to the user would change their perception of the implications of this technology. In *Fallible Digital Memory*, I challenged the faithfulness and certainty of a typical digital memory storage device by making a device that confuses and forgets. In *Tangible Rumination*, I speculated that encouraging the repetition of phrases and engaging with one's voice in a tangible way could make purposeless rumination into a form of interaction that encourages people to live in the present moment. It would be hard to imagine how a user who does not know the main

argument would approach a plastic bag with metal strips, a similar raincoat, or a vase-like structure on a wooded box. Therefore, the artifact itself does not hold all the knowledge.

There is then the question of how the knowledge could reach a bigger audience in the form of participation. Unfortunately, this question was not fully answered. I have been the primary user of all my systems. Self-usage had certainly been a source of knowledge that informed further development. For example, it was my own experience of the difficulties in training an ML system that prompted me to develop the argument of *Hidden Labor in ML*. It was also my own experience of testing the sound segmentation algorithm that inspired me to explore further the effect of repetition on one's attention to the meaning of thoughts. I echo the Background chapter that this autobiographical design approach could still be a valid approach to knowledge generation, as discussed in more detail in individual encounters.

Here I identified the sources of knowledge I learned about and how I attempted to contribute to the knowledge in these ways. To be able to claim that I had generated any knowledge at all, I need to evaluate my research against recognized standards. Zimmerman et al. (2007) proposed that Research-through-Design Projects be evaluated for invention, relevance, and extensibility.

I would like to claim that my work is inventive in many encounters. In *Intimate Ambiguity*, the innovation lay in the concept of designing ambiguous sensors, which was picked up in *Fallible Digital Memory* and *Tangible Rumination*. In *Hidden Labor in ML*, the innovation lay in the pipeline for training custom, on-body machine learning models. As for relevance, or how the work could push human knowledge or society forward, it is quite similar to the evaluative framework by Krogh and Koskinen (2020) for the programmatic tradition where good research should propose something new to the discussion. I believe that *Intimate Ambiguity* could prompt people to reflect on how technology relates to their body; *Hidden Labor in ML* brings to light the dialectics of training opaque, generic ML models with great effort versus training intimate models that are accessible to more people; *Fallible Digital Memory* and *Tangible Rumination* both speculate on how digital devices might shape alternative narratives around our mental processes. These four encounters are highly extensible as well. In many encounters, I discussed what did not work and what warranted shifts, but each of the obstacles could be an opportunity to develop further. For example, the ambiguous language generation could be a thesis in itself. Reinforcement learning and alternative musical instrument were potential topics. I could also

enforce machine forgetting algorithmically and stayed with *Fallible Digital Memory*, and I am sure the rumination mask could better support this activity in ways I did not have time to develop.

There are also many possibilities that could branch out at each of the drifting points, namely, what if I did not drift or had drifted in another direction? This proves that the Journey is more than where I landed by the thesis deadline, but rather a rich generative process that could be the starting point for further research on any concept this Journey touched upon and drifted through.

Discussion

Overview of the Discussion

I had no experience with design research when I first started this Master of Science in Computational Design (MSCD) program. Coming from a background in physics, I was intrigued to notice so many drifts that took me wandering in a vast landscape instead of what I thought would be a narrow road: as narrow as merely building an artifact involving sticky air pouches and finding a discourse around it. This surprise prompted me to document the research process carefully and eventually resulted in the bigger thesis wrapping around the research journey.

I would like to reflect on what is unique about design research and how it tolerates drifts more than what I was familiar with. To highlight this fact, the whole Journey chapter interweaves three kinds of text: letter, journal, and main text. The letter style features a continuous stream of narration about the research process. The journal part contains excerpts from my research journal that depicts the research experience at the moment. The main text was written in an academic tone and provides technical and literature information.

This interweaving of writing styles deviates from how a traditional research publication is often written to feature a linear path toward success and challenges the definition of a good piece of research writing: writing that conveys knowledge.

Writing this thesis was an autobiographical design process as I wrote, edited, and adjusted the delivery of the information. The outcome is a new design method to design how the diverse forms of knowledge in design research could be presented. This design method acknowledges and celebrates the drifts, and brings the reader into the ride where knowledge is produced.

Upon interviewing five students in the MSCD program, I found a majority of them voicing that more documentation of the drifts would help the readers to understand and follow the research pro-

cess, though one student argued that more documentation could bring the readers further from the core concept the thesis was designed for. With the evaluation from four MSCD students, my documentation method is effective in bringing the reader to my feelings and thoughts, and the incorporation of the letter style provides a lighthearted, easier read.

It is my humble wish that this thesis could serve as a starting point for future students to expect and embrace the drifts; for faculty members to encourage exploration of their design research space; and for opening up the discussion in the design research community around what interdisciplinary inquiry indicates in the documentation process for better communication of knowledge.

Drifting from Physics

During the writing of the Journey chapter, I summarize the forces that caused drifting in my research into the following categories:

- *Internal forces*: my personal interests, my existing skills, and my reasoning of what arguments are coherent and meaningful.
  - Example: When I reasoned that the mapping from human emotions to body gestures did not make sense, I drifted away from Intimate Ambiguity.
- *Artifact affordances*: what worked and did not work during implementation and self-usage.
  - Example: When the perceived effort in interacting with my ML system did not align with the argument of bringing awareness to the hidden labor in ML, I drifted towards implementing new sound processing techniques and landed in Fallible Digital Memory.
- *External information*: advisor and peer feedback and literature review.
  - Example: When S. Goral pointed out that ambiguity was an elusive concept and the interaction with machines could be a rich design space, I moved to designing the ML control system.

These are probably hard to find in the physics research I had been familiar with, or at least not to the same extent. In my three years working as a research assistant in a physics lab that focuses on high-energy beam physics and plasma physics, I was mostly working with

computer simulations and occasionally with the fabrication of nanostructures. The workflow I perceived in their way of doing research was that we would start with a need, for example, "We want to know the energy spectrum of a gamma beam". This need would serve as the constraint we design within. After the need was identified, we would go on to invest time and effort into designing a system that could measure the energy spectrum, for example. The researcher, who provided me with simulated data, changed the design of his device in the simulation at least five times, which resulted in me changing the data analysis pipeline significantly for multiple rounds. However, the need for measuring the gamma-ray spectrum did not change. The result was that the researcher had a better and better design of the simulation, and my data analysis pipeline improved as well.

The above example illustrates that I am more familiar with the way that the "argument", or goal in this case, influences how the artifact changes and not the other way around. If the simulated device had a superpower in measuring photons of lower energy, it would not make it useful to support the goal. In the words of Zimmerman et al. (2007), this is a very finely constrained problem, more commonly seen in engineering.

Drifting Categories by Krogh and Koskinen (2020)

Therefore, I sought confirmations from literature that describes the drifts in design research. Krogh and Koskinen (2020) summarized four ways drifting might happen in the process of generating knowledge in design research, which they call epistemic traditions.

- Experiential: modifying the artifact to allow it to produce the message the researcher wants to convey.
- Methodic: changing the method of investigation based on explicit evidence.
- Programmatic: modifying the scaffolding frameworks and theories that guide the design research.
- Dialectic: shifting the research direction based on feedback from the users of the research artifact or through a conversation with users.

It is not fair to force the drifting points I summarized from my experience into the four ways of drifting presented by Krogh and Koskinen (2020), since the drifting points I summarized are categorized based on the source of influence, while their analysis focuses

on where drifting occurs in the design research processes. However, I could identify some interrelationships between these two ways of analyzing drifting.

The internal forces I identified could drive the drifts in all epistemic traditions. My personal interests could change the message I want to convey in the artifact (experiential) and the frameworks I would like to design in (programmatic). My existing skills would change my method of investigation (methodic). My reasoning of what arguments I would select could influence the further development of the artifact (experiential) as well as the direct framework and theory I deem fit (programmatic).

The artifact's affordances mostly drive the drifts in experiential and programmatic traditions. In other words, what worked goes into the artifact and the argument, and what did not work changes the artifact and the argument.

External information would be characterized as applying to the dialectic tradition, however, they also played an important role in the early stages such as forming the framework (programmatic) and changing the artifact to convey the framework (experiential).

If we take an even closer look at the drifting points that I summarized above, one source of influence often results in an imbalance of another force, thus pushing the research forward in zig-zag paths. One may argue that this is how design research happens. In particular, I found it very common (and hopefully evident in the previous chapter, too) that programmatic drift would result in experiential drifts, and vice versa. For example, when I decided to explore the human-machine tension in intelligent systems instead of the ambiguity of intimacy, I quickly shifted my effort into designing a machine learning system that forgot information. On the other hand, when I noticed that repeating phrases over and over again did not make me forget, but rather stripped off their meaning, I drifted my argument from allowing digital devices to forget to allowing them to ruminate with us without purpose.

Designing this Thesis

By now I have discussed how I drifted in the research process and how they relate well to the observed drifting points in design research, and I would now like to discuss the method of addressing these drifts in a novel way, and how it improves upon the documentation methods in the Background.

As I stated in the Introduction, "this thesis is about the process of

constructing design research theses rather than a collection of unfinished theses I have made". Making a thesis shares a lot of common methods and insights with making a design research artifact. It is a form of knowledge production through the process of gathering materials, forming arguments, and organizing arguments into reasonable structures. The craft aspect of this artifact could even be as concrete as typing the words on my computer and preparing the figures. I think this parallel is helpful for understanding that the documentation of the four encounters is the "methods" section of this bigger thesis.

When I treat the encounters as design artifacts and study how I made them, I am using autobiographical design to design this thesis. This thesis meets the definition of autobiographical design research which requires the researcher's genuine needs and long-term use of the researcher (Neustaedter & Sengers, 2012). My genuine needs, as stated in the Introduction, were not only my needs to have a thesis to graduate but also the passion to explore myself as a new researcher and how I may contribute to the discourse around design research. The way I "used" and "lived with" this thesis is embedded in the development, documentation, and hours of pondering over research arguments in the past months. For example, I have a three-thousand-word journal that I kept since Janurary, numerous presentation slides in courses, and documents or proposals of every encounter I discussed from which I extracted the forecasts and imaginary introductions.

Types of Writing

The most dominant feature of this thesis is the weaving structure of three styles of writing: letter, journal, and main text.

The letter-style writing was inspired by Jack Kerouac's *On the Road* (1999), through which he presented his definition of "spontaneous prose". To digest his writing style, I interpret his work as a stream of narration featuring event descriptions in chronological order and various side comments and reflections. The letter style of the Journey chapter was written in a sprint of four hours followed by careful editing that filled in missing gaps of information and adjusting of the tone of voice. It is casual, and reflective in terms of incorporating inner dialogues, and I intended it to sound like a letter addressed to a friend. I think weaving this prose-like writing with formal writing in this thesis is valuable since it truly reveals the messy process of developing a thesis in design research and brings the reader closer to what is commonly hidden or edited out of formal publications.

For example, this excerpt addresses the reader as a letter often does, and refers back to the previous discussion of how music performance was not well received in the beginning. I then move on to discuss the implementation of the system (p. 67):

> Let us go back to the concept of music performance, although it was already born with ominous signs. For ease of testing, I also use musical notes as the output of the model, so I would know if the model is predicting the same note that I sang.

In general, the letter style captures my reflections on the journey, as well as an overview of the thought process and decision-making during the journey.

The journal entries are excerpts from my research experience documented at the moment (p. 68):

> *[1/23] I added sound after I saw that the machine could produce 3 individual segments without my guidance, because I wanted to know if they were actually accurate.*

There are often missing words, grammatical errors, and illogical descriptions. However, I argue that those are evidence of how I worked as a researcher in this messy process.

The main text is standard academic writing, typically found in any research publication (p. 52):

> This thesis addresses this question by developing and testing an on-body system of air pouches that serves as sensors of body gestures, and an associated natural language model to generate language about intimacy. Each set of similar body gestures is mapped to elements in the language model, which is customarily trained to produce ambiguous output in the form of poetry.

There are both *Landing* sections and *Discussion* sections for each encounter. The Landing sections read similar to a standard Introduction to a paper. They help ground the readers in the concepts the journey is operating in.

Weaving

I weave these three styles together to achieve various effects. They are similar to musical instruments that compose a symphony, as they elicit different thoughts and reading experiences while the whole is more than their individual abilities.

Example 1 The letter provides personal motivations, and the main text provides background literature to enrich the research (p. 42).

> Looking back from where I am now, there were definitely personal motivations that shaped the thesis into what it is now. First of them is my fascination with the human body. From the pouches, I decided, a somewhat forceful stretch, that I wanted to put it on someone's body. Maybe it was the tactile feeling of a soft material, maybe it was the inflation that was reminiscent of a fish's sac.

This thesis is inspired by the concept of somaesthetic design, which hypothesizes that awareness of the world and potential improvement in the ways we live and act could be achieved through body awareness. (Höök et al., 2015). Somaesthetic design is related to embodied interaction design as they both consider tangibility as important, but somaesthetic design emphasizes designing for experiences (Höök, 2018). It is also different from affective computing where feelings are studied as objects, and instead includes the felt experience in the world. As for aesthetics, somaesthetic experiences are completed in dialogues between the users and the design.

Example 2 This example interrupts the flow of the letter to provide technical details of the implementation (p. 46).

> Ambiguity came into play for this reason as well. I first characterized the performance of capacitive sensing based on the Multi-Touch Kit work,

L. Yao suggested that I consider capacitive sensing as a way to digitally capture the shape transformations of the pouches. With Professor Daragh Byrne's suggestion of manufacturing custom capacitive sensing structures with the Multi-Touch Kit (MTK) (Pourjafarian et al., 2019), I downloaded the code and built the structure as they specified.

> from which I learned that the Arduino analogRead signal was linearly proportional to the capacitance of the sensing electrodes I attached it to.

Next, I measured the signals from a single MTK electrode pair and came to the conclusion that the analog signals from the Arduino were roughly proportional to the capacitance between the electrode pair.

Example 3 The second paragraph foreshadows future development and links the letter together, making it more reflective (p. 56).

Since the mapping from emotion to gestures and the natural language generation for ambiguity was too open to tackle, I made the above mind map to think through what else I could do with machine learning. Fig. 21 classified Machine Learning (ML) in terms of its input and output and arrived at the Explainable Artificial Intelligence (AI) and Tangible AI concepts.

> What I did not anticipate was that the sensing mechanism eventually revived the concept of ambiguity, or uncertainty, or any of that fuzzy ball of concepts at later encounters of the development.

Example 4 Three styles provide different information on the same event of singing. The letter provides an overview of the experience: "richness". The main text describes what I did to test, and the feeling of "surprise". The journal provides the most intimate experience: I felt "self-conscious" and "embarassed" (p. 73).

> My roommate said she heard me singing all night. This kind of labor was again unexpected. The moment I heard the mixed sound played back to me like a much delayed but everlasting echo, I was amazed by the richness of associations I had.

The first couple of tests were me humming some short melody, so the playback was simply like a chorus of multiple versions of me. Then I started testing with word phrases, which the system was also able to record and mix. Having phrases of varying loudness and lengths looping in a mix, was a very surprising experience for me.

> *[2/7] [...] which makes me extremely self-conscious and embarrassed (?) by the sound I repeat. Therefore I prefer humming to words, but also short complaints and humorous remarks about my state of life.*

Example 5 Sometimes, when read together, the main text displays how the struggles in the letter were (temporarily) resolved. Here the struggle is deciding whether the phrases should be negative (p. 88).

> Again, I started to struggle with framing the argument: are the phrases being repeated bad or negative? Or can they be neutral?

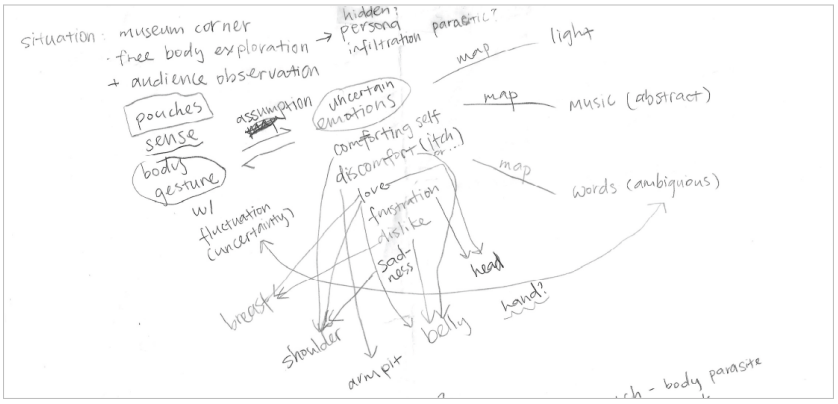
[Landing: Tangible Rumination] When we think of something repeatedly, we might be "ruminating". Merriam-Webster's Dictionary defines "to ruminate" as "to go over in mind repeatedly and often casually or slowly" (Merriam-Webster, n.d.). This seemingly counterproductive activity also "suggests little of either purposive thinking or rapt absorption". Rumination in our brains contrasts with software and devices which can recall and present to us our recorded thoughts in a linear fashion as efficiently as a click of a button. However, while productivity is valuable, purposelessness has its unique role in life as well, such as bringing back playfulness and breaking the momentum of analyzing life. This thesis examines the potential for digital devices to mediate rumination instead of eliminating it.

Formatting and Interstitials

As displayed in the examples above, the formats of the three types of narration are different. The letter occupies a right-shifted block and is in a lighter shade of gray. This allows the reader to read the letter as a continuous stream of narration. The journal entries are centered and are displayed in a font similar to that on a typewriter, enhancing the association with a traditional research log. The main text occupies the original text column in the Overleaf Tufte-Style Book Template [12] that I am gratefully using.

There is a wide margin on the right of the main text column, where I put some of the less essential images. Most images take the main column, and sometimes the full width.

The journey includes many scanned hand-written notes, for example in Fig. 19 and repeated here. These scanned notes often capture



brainstorming activities and transitions that resulted from these extensive reasonings.

However, there are other images, unlabeled and without captions, that are called interstitials. Their introduction was inspired by W. G. Sebals's Austerlitz (2011), which included many enticing and mysterious photographs in the middle of his narration. My interstitials were drawn on paper and scanned as well, and they do not contain specific information related to the text or the journey. Fig. 55 and Fig. 56 are two examples of interstitials.

These unlabelled images act as breathing points for the readers to gather their thoughts and to freely associate the abstract pattern with the content.

[12] Kleb, B., Wood, B., & Godby, K. (2017). Tufte-latex. Retrieved April 22, 2023, from https://www.overleaf.com/latex/templates/a-tufte-style-book-with-vdqi-title-and-contents-page/xqkjxvmrhmmp



Figure 55: A shorter (and therefore bigger) interstitial.

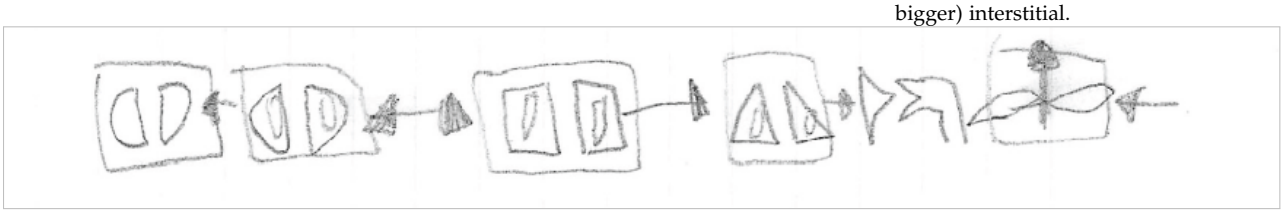


Figure 56: A longer interstitial.

Evaluating this Thesis

This thesis is highly experimental. In the Introduction, I outlined several criteria for evaluating the writing of this thesis: I would like the thesis to be critical of the attempted rigor at each encounter, reflective of motivations and forces that prompted drifts, intimate between the reader and my inner thoughts, explorative in terms of narrative styles, and receptive to alternative ways to approach thesis work.

Neustaedter and Sengers (2012) equate rigor with critical reflection of the work processes in autobiographical design and quote from Greenberg and Buxton (2008): "At a minimum, authors should critique the design: why things were done, what they learned, expected problems, how it fits in the broader context of both prior art and situated context, what is to be done next, and so on.".

Therefore, I discuss the first two criteria here. As I wrote the Journey, I reflected on the rationale behind my research choices. What I learned from the (un)expected problems during development often prompted drifts that changed how things were done. Throughout the Journey, I included background literature that situated each of the encounters in the broader context. In the discussion of each encounter, I proposed future work as well.

To support this evaluation, I held conversations with students working on design research theses in the MSCD program or graduates who had finished their theses in the previous year. I did this in order to investigate if the drifting process illustrated here resonates with other students' experiences, especially those coming from non-technical fields. This is the receptive part of this thesis. Chengzhi

Zhang and Tomas Cabezon Pedroso are working on the final stages of their design research theses. Zhenfang Chen, Yanwen Dong, and Malika Khurana have completed their MSCD theses in May 2022 (Chen, 2022; Dong, 2022; Khurana, 2022). All of them had experience with design in the past, though the degree they consider their experience to be research varied.

Drifts and the Practice of Documenting Drifts

From the conversation, I could identify most sources of drifts that I experienced in these students' experiences as well. For C. Zhang, after developing a speculative design artifact, she realized that it did not align with what she would like to do in the future, so she moved away from that topic. This drift is similar to the moment I turned away from natural language generation when I realized that it would be so much work that it would become another thesis. I consider this kind of drift to be experiential since it originates from the artifact. Other people's feedback, or along the line of dialectic tradition, greatly influenced Z. Chen's development and forced him to refine and scope his thesis in the early stages. Emek Erdolu, a PhD student in our program, posed questions about Z. Chen's motivation to work on the research, its significance, other related work, and so on. Later, when Z. Chen was working with another more developed idea, his advisors again challenged him with the relevance of the project to his research argument, and prompted him to drift to his final artifact. Skills and implementations, the source I identified in my experience, limited C. Zhang to push her original goal of developing a full-fledged AI platform for designers into her future plan, because she still needs to learn more coding skills and stay informed about the evolving practice. I had the same experience with the RL agent, where my limited knowledge could not support the implementation and I drifted to a passive classification task.

An interesting source of drifts that I did not include above in the discussion was the time constraints. For Z. Chen, he struggled a lot in January and Feburary of 2022.

> When it goes to the middle and end of February, I'm pretty worried because I cannot submit anything to Daniel[13].

[13] Our academic advisor Daniel Cardoso Llach.

Eventually, he forced one artifact out of his previous project, and felt somewhat disappointed. Apart from the lack of skills, C. Zhang also acknowledged that the time constraint was a major factor for her to scope down the thesis. Y. Dong expressed that her original plan was to do a lot more prototypes at larger scales, but the time constraint

only allowed her to finish two smaller projects and a proposal for the third.

Despite the prevalence of drifts, there are notable obstacles towards documenting them. One practical reason T. Cabezon Pedroso and Y. Dong touched upon was that when things don't work, they are deleted. Code, for example, was hard to keep track of for T. Cabezon Pedroso because when it was not under version control, the failures disappear completely. This makes it hard to recover the design decisions made at the moment. The result was that most of the drifts were not documented in as much detail as I have done here. For Z. Chen, he did not put the earliest idea in the final document, but included parts of the later experiments in a sequence and tried to make a storyline for what he did. For parts he omitted, he felt like they were not as relevant as the project progressed. Y. Dong also did not include much of the drifts, and "worked around" them to land on the two successful experiments.

However, all of my interviewees expressed that drifts are worth documenting. T. Cabezon Pedroso, Z. Chen, and Y. Dong consider it so because documenting drifts could help future researchers or people who want to follow what they did to avoid the dead ends and pitfalls they have encountered. Y. Dong stated:

> I think it is part of the comparison between like, your design, thoughts and process versus the actual execution and this should be part of the documentation.

T. Cabezon Pedroso gave a higher-level argument, saying that society, as evidenced by social media, tends to show the "good things", while the other side is also needed. M. Khurana also stated:

> I think it's great that you're talking about this, because in master's theses and academic papers, it feels like there's a lot of pretending that happens, for you to act like it was a linear path of like this was the technology that we were interested in using, this was the idea, and this is what we made, and this is how we tested it.
>
> But the more I talked to people it seems like it's such a mess like there's a moment in every single paper where everyone is like what the hell did we do? Like how do we make sense of this? And yeah, it frustrates me that that's hidden a lot of the time.

It is also hard to find traces of drifts in published design research articles. Many interviewees provided similar reasons, stating that for papers, especially conference papers, the results matter more.

Y. Dong noted that formal publications focus more on results, and the process that leads to successful results, so that people could follow the right path. C. Zhang argued that:

> For the conference papers, I think it's not necessary to present the process, because you're only conveying your research findings. So the process does not matter. If you're presenting a whole design project, it's not. It's not just about the final rendering or the final product, which is just the product itself. So there's right. Yeah. Like, the process also matters.

When and How to Document Drifts

In this thesis so far, I have had the assumption that people do not document drifts enough. However, many interviewees pointed out that documenting drifts is quite common in design practices, and in theses related to design. C. Zhang showed me her documentation of a few projects on her website. She believed that by showing the drifts in the iterations, reviewers and readers could be more convinced that the researcher/designer has thoroughly considered the alternatives and arrived at the optimal solution. This view echoed with Neustaedter and Sengers (2012), who believe that a true design process could help "reviewers to better judge the work".

T. Cabezon Pedroso and M. Khurana described that, if they were to document drifts, they will be included in the Discussion. T. Cabezon Pedroso reasoned that because he wanted to treat his three experiments as independent and publishable conference papers,

> I realized that because of that, there is not problems. So then on my discussion side, I'm discussing all the things that are behind the scenes. So I'm actually saying, oh, there is all these behind the scenes that happened in research that hasn't been shown.

M. Khurana gave a more detailed reasoning. Although she kept reminding herself to be honest, she only talked about the drifts to some extent. She recognized that writing about the process may not be as engaging for the reader. More importantly,

> to go into detail just becomes about that, then it's not about the concept or the project. You got the process and what is the process of coming up with creative ideas and following through on them. And I didn't want to write a thesis, that was just about that process. I wanted to be about the concepts.
>
> I think it can dilute the point of the thesis.
>
> I think some people might think that it delegitimizes it, I feel like, it makes it less clear what the thesis is about.

To balance talking about drifts with the concept, she proposed:

> In a thesis, I especially because I wanted it to still be about the projects themselves, I would want to, like, start off very clear, like, this is the theme and these are the projects I'm writing about, to give the reader a very concrete idea of what you did. Yeah. And then the discussion and conclusion are where you can, like, make a mess again. You've communicated to them clearly what, what you're trying to do, and then you can be like, okay, but here's like all the complications with it.

The Effectiveness of My Approach

I selected four participants, Yanwen Dong (graduated), Zhenfang Chen (graduated), Malika Khurana (graduated), and Tomas Cabezon Pedroso (in progress), and provided them three pages from the compiled Journey along with three followup questions. The three pages contain all three types of writing, a provocative interstitial, and a figure. The details of the pages and the questions can be found in Appendix C: Interview Questions.

Overall, I received positive feedback regarding the readability of the text. For Y. Dong, she stated that "It was confusing at first glance", but she could follow after half of the first page. T. Cabezon Pedroso stated that he re-read the document after reading the styles consecutively, but only the letter. According to Z. Chen, the different formatting helps with signaling the transition of tone and styles. M. Khurana expected the styles to be "in direct dialogue with each other", but that turned out not to be the case.

In terms of conveying the research process, my approach brought them into an atypical reading experience. T. Cabezon Pedroso commented that "it is like reading a book not like reading a thesis". He also recognized that the letter style conveys the feelings, frustrations, and the "space for error" that are often missing in research documents. Similarly, Y. Dong read more "self-reflections and critique" from the letter style, and found it more lively and descriptive "as if you are trying to make the reader experience the process through words". Z. Chen mentioned that the interweaving of the letter and the main text is interesting because the main text reads more serious and rational, which could take some time for him to digest, while noticing that the next paragraph is the letter could be a "buffer point" that makes him more relaxed.

Only Z. Chen noticed the interstitial image and mistook it for a draft image of my prototype. This was somewhat intentional, as with Sebald's image. There is no mention of how this affected the read other than making it "lighter".

As for alternative ways of documenting the drifts, M. Khurana would consider using "a flow diagram or timeline of the process"

to document this drift. Y. Dong envisioned two forms that could be more natural:

> To me it would flow more naturally if you write in full "letter style" as a piece of design research writing, or in a combination of "academic style" adding "research journal" to show contrast between prepping and what actually happened?

These are valuable suggestions as they deviate less from the traditional documentation, and could be a fallback if and when my approach becomes less ideal in situations.

Conclusions

In conclusion, this thesis revealed the messy process under which design research generates knowledge beyond the traditional scientific and engineering methods. The effort to develop and frame the core artifact, an interactive air pouch with a sticky interior, was full of drifts, defined as moments design research changes the question or focus and lands on something unexpected. In this journey, I drifted through four major encounters with different discourse around the artifact, which were *Intimate Ambiguity*, *Hidden Labor in ML*, *Fallible Digital Memory*, and *Tangible Rumination*.

I use three styles of writing to convey the process of drifting. The letter style addresses the reader in a continuous form of narration, and provides insights and reflections on the research process. The journal entries documents the experiences and decisions at different moments during the development. The main text uses an academic tone to supply technical information and background literature. The dense weaving of these three styles and the insertions of scratch paper sketches faithfully capture the multi-dimensional and multi-agent process of design research. Through this thesis, I have proposed that this documentation method could celebrate the drifts as often hidden in design research publications and bring the readers to the core of how knowledge is produced. Furthermore, this work invites the design research community to innovate on documentation methods and other creative ways to disseminate knowledge.

Limitations and Future Work

I am aware that these findings and evaluations are highly personal. As with autobiographical design, the findings generated are only experiential understandings and cannot always be generalized (Neustaedter & Sengers, 2012). If another researcher were given the starting air pouches, they would most likely not make any of the four encounters into a full thesis, and more importantly, they would probably not arrive at the reflections and documentation of drifts as this thesis is now about.

As evident in the conversations above, there is no right or wrong way to present drifts, and it often relies on the focus and goal of the research document. Logic and clarity often demand structures that the drifts defy. Desjardins and Key (2020) use diverse symbols to illustrate the design process in their paper, yet they still attempt to distill and abstract the process with labels. I believe that this is a major trade-off in this line of work between clarity in writing and rationality in discussion versus truthfully capturing the deep weeds researchers wander into, risking the work being incomprehensible. My attempt to weave in prose-like writing was to address this trade-off.

There is also a trade-off between the ability to maintain a document as a team versus documenting the internal processes as an individual researcher. The reflections and feelings may not be easily shared in hindsight, as two researchers writing about their journey could feature different lines and temporal order, making the interweaving difficult.

As evident in the retrospective letter, my thesis is a device to communicate knowledge production and its process, rather than a document that supports the design research in the moment. Many of the forms of documentation in the Background are supporting the design. I argue that my format can be considered as harvesting the supports from the process to present them in the final document.

The Background chapter described documentation as a research process. I certainly believe that this is an area where innovation

should be encouraged. I would like to invite other students in design research to expect and embrace the inevitable drifts. I have received great support from my advisors and other faculty members, and I would like to see how they could pass the encouragement to experiment with writing further.

Meanwhile, I am aware that written and graphic materials pose extreme limitations on what and how research can be documented, and I would like to argue for better access to multimedia information and retrieval so that research is no longer bound to a linear narration or a 2D plane. For example, the MediaObject project by Sadokierski (2019) publishes a hybrid of audio-visual documentation of design processes/performances, and interviews, besides the traditional text-based document. Their publications are also double-blind peer-reviewed. Another example of multimedia publishing is *Scalar* by Alliance for Networking Visual Culture (2021), which features non-linear lines of narration and annotations of digital objects with an interface similar to a blog page. I believe that this could be the future of design research publications.

From the perspective of a student with a technical background learning to conduct design research, I wish to contribute one case to the conversation of how knowledge from design research could be conveyed. I believe that my unique perspective would be valuable for future interdisciplinary research, where various research methods and documentation methods clash and mix to produce unusual works, one of which this thesis is striving to become.

Acknowledgements

Here I would like to express my sincere gratitude towards the people who supported me in the journey that has been full of drifts. Many of you directly or indirectly pulled this project in stray directions which later proved to be the direction forward.

First of all, I would like to thank Professor Daragh Byrne for his understanding and encouragement during the two years of ups and downs. Your surveying knowledge of the field has become the comprehensive map that grounded my research. You have also revealed to me the charm of interdisciplinary research situated at the crossroads of the technical and the designerly and have rightfully rejected this dichotomy.

I would like to thank Professor Lining Yao for sowing the seed of tangible interaction that has been the core of this thesis. The project that originated from your class has led me through iterations of explorations in this for me uncharted territory. Here I would also like to acknowledge the valuable inspiration and collaboration I have received from my classmates Violet Han and Di Wu during the development of the very first sticky pouches.

Sinan Goral has brought me burning questions and fueled my critical interrogation of the core argument my thesis now supports. Your comments of this work as "great and weird [stuff]" prompted me to put in "a stupid amount of effort" to kindle the wildest fire I have ever dared to ignite across the pages you now have in your hands.

Next, I would like to thank Professor Marti Louw for the suggestion of using Edward Lear's works as potential outputs of my machine learning system, and Professor Jesse Stiles for showing me how to use Max MSP to manipulate sounds. Thank you Professor Alexandra Ion for the earliest explorations.

Many thanks to all the friends and colleagues who have been there throughout the journey as well. To Computational Design alumni Zhenfang Chen, Yanwen Dong, and Malika Khurana who have pro-

vided so much feedback and mentorship. To my research mentor Jingyang Liu who I envisioned and sometimes fantasized with about the potential of this thesis. Thanks to my colleague Tomas Cabezon Pedroso for the friendly discussions and the fun we had while struggling together. To Chengzhi Zhang for her insight into design practice and research. To Howard Han for the mutual support and encouragement as well as the hilarious moments. To Hongyu Mao for being my work buddy. To Suzie Liu, Jingxuan Guo, and Chelsea Lai for supportively listening to my ever-drifting framing and arguments.

This thesis would not have been possible without the support of my mother, Yanli Xie, who spent nights during my lost mornings and mornings during my anxious nights analyzing my hypothesis, questions, and arguments with me, often taking notes of her own.

The core materials of this thesis largely come from the generosity of Professor Olivia Robinson and Dr. Garth Zeglin, or were purchased with the Graduate Small project Help (GuSH) Research Grant.

The formatting of this thesis uses the beautiful code in the Overleaf template by *The Tufte-LaTeX Developers* at https://www.overleaf.com/latex/templates/a-tufte-style-book-with-vdqi-title-and-contents-page/xqkjxvmrhmmp.

Lastly, being a member of the MSCD program and of the web of relationships at large, I find it impossible to enumerate the uncountable noun of gratitude. To all of you whose life trajectories intersected mine, I will carry your support forward and extend it to those who may need them as I continue to drift in this world.

Bibliography

Agrawal, H., Yamaoka, J., & Kakehi, Y. (2018). (author)rise: Artificial Intelligence Output Via the Human Body. *Proceedings of the 23rd International Conference on Intelligent User Interfaces Companion*, 1–2. https://doi.org/10.1145/3180308.3180328

Agre, P. E. (2014). Toward a Critical Technical Practice: Lessons Learned in Trying to Reform AI. In *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide* (pp. 131–157). Psychology Press. Retrieved April 19, 2023, from https://books.google.com/books?hl=en&lr=&id=XSaPAwAAQBAJ&oi=fnd&pg=PA131&ots=odgXY-zV10&sig=ABtpp58widanxlb-k26r82qGp4s#v=onepage&q&f=false

Alliance for Networking Visual Culture. (2021). Scalar. Retrieved May 8, 2023, from https://scalar.me/anvc/scalar/

Bae, J., Cho, H., Song, K., Lee, H., & Yoo, H.-J. (2012). The Signal Transmission Mechanism on the Surface of Human Body for Body Channel Communication. *IEEE Transactions on Microwave Theory and Techniques*, *60*(3), 582–593. https://doi.org/10.1109/TMTT.2011.2178857

Bardzell, J., Bardzell, S., Dalsgaard, P., Gross, S., & Halskov, K. (2016). Documenting the Research Through Design Process. *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 96–107. https://doi.org/10.1145/2901790.2901859

Baumer, E. P., Khovanskaya, V., Matthews, M., Reynolds, L., Schwanda Sosik, V., & Gay, G. (2014). Reviewing reflection: On the use of reflection in interactive system design. *Proceedings of the 2014 conference on Designing interactive systems*, 93–102. https://doi.org/10.1145/2598510.2598598

Bere, C. (2002). The book of memory: W.G. Sebald's The Emigrants and Austerlitz [184]. *The Literary Review*, *46*(1), 184+. Retrieved April 24, 2023, from https://link.gale.com/apps/doc/A94983806/AONE?u=cmu_main&sid=googleScholar&xid=d0e2e783

Bertran, I. (2014). Memory Device. Retrieved April 21, 2023, from http://ishback.com/memory/index.html

Birch, S. (2013). Freakshow. Retrieved April 21, 2023, from https://ohbythewayblog.blogspot.com/2013/08/beauty-mixed-media-and-photography.html

Boem, A., & Troiano, G. M. (2019). Non-Rigid HCI: A Review of Deformable Interfaces and Input. *Proceedings of the 2019 on Designing Interactive Systems Conference*, 885–906. https://doi.org/10.1145/3322276.3322347

Boucher, A., Gaver, W. W., Pennington, S., & Walker, B. (2001). *Equator workbook one: Ideas, scenarios and proposals for the home.* Royal College of Art.

Bourgault, S., Forgues, E., & Jacobs, J. (2020). Experimenting with Robotic Softness [event-place: Montreal QC, Canada], 71–75. http://www.isea-archives.org/docs/2020/ISEA2020-Proceedings.pdf

Carroll, J. M., & Kellogg, W. A. (1989). Artifact as theory-nexus: Hermeneutics meets theory-based design. *Proceedings of the SIGCHI conference on Human factors in computing systems Wings for the mind - CHI '89*, 7–14. https://doi.org/10.1145/67449.67452

Chen, Z. (2022). *Just Listen: Prototyping with TinyML to Augment Everyday Sound* (Doctoral dissertation) [Artwork Size: 11208410 Bytes Publisher: Carnegie Mellon University]. https://doi.org/https://doi.org/10.1184/R1/21208226.v1

Dalsgaard, P., & Halskov, K. (2012). Reflective design documentation. *Proceedings of the Designing Interactive Systems Conference*, 428–437. https://doi.org/10.1145/2317956.2318020

Demir, A. D., Kuusk, K., & Nimkulrat, N. (2022). Squeaky/Pain: Articulating the Felt Experience of Pain for Somaesthetic Interactions. *Temes de Disseny*, (38), 162–178. https://doi.org/10.46467/TdD38.2022.162-178

Desjardins, A., & Key, C. (2020). Parallels, Tangents, and Loops: Reflections on the 'Through' Part of RtD. *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, 2133–2147. https://doi.org/10.1145/3357236.3395586

Desjardins, A., & Wakkary, R. (2016). Living In A Prototype: A Reconfigured Space. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 5274–5285. https://doi.org/10.1145/2858036.2858261

Dong, Y. (2022). *Designing with Senses: Creating Interactive Multisensory Experiences in Public Space* (Doctoral dissertation) [Artwork Size: 10181085 Bytes Publisher: Carnegie Mellon University]. https://doi.org/https://doi.org/10.1184/R1/21331158.v1

Dunne, A., & Raby, F. (2013). *Speculative everything: Design, fiction, and social dreaming*. The MIT Press.

Dunne & Raby. (1995). Faraday Chair. Retrieved April 19, 2023, from https://collections.vam.ac.uk/item/O63805/faraday-chair-chair-dunne-raby/faraday-chair-chair-dunne--raby/

Eikey, E. V., Caldeira, C. M., Figueiredo, M. C., Chen, Y., Borelli, J. L., Mazmanian, M., & Zheng, K. (2021). Beyond self-reflection: Introducing the concept of rumination in personal informatics. *Personal and Ubiquitous Computing*, *25*(3), 601–616. https://doi.org/10.1007/s00779-021-01573-w

Eriksson, S. (2011). Intangible Embrace. Retrieved April 21, 2023, from https://www.konstfack2011.se/bachelor/baam/sofia-eriksson/

Eshel, A. (2003). Against the Power of Time: The Poetics of Suspension in W. G. Sebald's "Austerlitz". *New German Critique*, (88), 71. https://doi.org/10.2307/3211159

Fang, L. (2013). Liu Fang SS13: Sculpted Cashmere [photo credit: Quinn Connors]. Retrieved April 21, 2023, from http://pretemoiparis.com/2012/10/liu-fang-ss13/

Frayling, C. (1994). Research in Art and Design. *Royal College of Art Research Papers*, *1*(1), 1–5. Retrieved April 24, 2023, from https://researchonline.rca.ac.uk/384/3/frayling_research_in_art_and_design_1993.pdf

Fu, Y., Byrne, D., & Shea, L. (2021). Evoking the Post-industrial Landscape Memories through Spectrality and Mixed Reality Soundscapes. *Creativity and Cognition*, 1–6. https://doi.org/10.1145/3450741.3465256

Gaver, B., & Bowers, J. (2012). Annotated portfolios. *Interactions*, *19*(4), 40–49. https://doi.org/10.1145/2212877.2212889

Gaver, B., Dunne, T., & Pacenti, E. (1999). Design: Cultural probes. *Interactions*, *6*(1), 21–29. https://doi.org/10.1145/291224.291235

Gaver, W. (2011). Making spaces: How design workbooks work. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1551–1560. https://doi.org/10.1145/1978942.1979169

Gaver, W. W., Beaver, J., & Benford, S. (2003). Ambiguity as a resource for design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 233–240. https://doi.org/10.1145/642611.642653

Ghajargar, M., Bardzell, J., Renner, A. S., Krogh, P. G., Höök, K., Cuartielles, D., Boer, L., & Wiberg, M. (2021). From "Explainable AI" to "Graspable AI". *Proceedings of the Fifteenth*

International Conference on Tangible, Embedded, and Embodied Interaction, 1–4. https://doi.org/10.1145/3430524.3442704

Green, K. E. (2019). Strange Places: Loci of Design Inspiration. *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction*, 173–180. https://doi.org/10.1145/3294109.3300998

Greenberg, S., & Buxton, B. (2008). Usability evaluation considered harmful (some of the time). *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 111–120. https://doi.org/10.1145/1357054.1357074

Hitsuwari, J., Ueda, Y., Yun, W., & Nomura, M. (2023). Does human–AI collaboration lead to more creative art? Aesthetic evaluation of human-made and AI-generated haiku poetry. *Computers in Human Behavior*, *139*, 107502. https://doi.org/10.1016/j.chb.2022.107502

Honnet, C., Perner-Wilson, H., Teyssier, M., Fruchard, B., Steimle, J., Baptista, A. C., & Strohmeier, P. (2020). PolySense: Augmenting Textiles with Electrical Functionality using In-Situ Polymerization. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3313831.3376841

Höök, K. (2018). *Designing with the Body: Somaesthetic Interaction Design*. The MIT Press. https://doi.org/10.7551/mitpress/11481.001.0001

Höök, K., Ståhl, A., Jonsson, M., Mercurio, J., Karlsson, A., & Johnson, E.-C. B. (2015). COVER STORYSomaesthetic design. *Interactions*, *22*(4), 26–33. https://doi.org/10.1145/2770888

Jang, S. (2013). Calibre. Retrieved April 19, 2023, from https://sungjang.com/work/calibre1-gy8k2

Johansen, S. S., van Berkel, N., & Fritsch, J. (2022). Characterising Soundscape Research in Human-Computer Interaction. *Designing Interactive Systems Conference*, 1394–1417. https://doi.org/10.1145/3532106.3533458

Karami, S. (2018). *Interruption: Writing a Dissident Architecture* (Doctoral Thesis). KTH, School of Architecture and the Built Environment. Retrieved May 6, 2023, from http://kth.diva-portal.org/smash/record.jsf?pid=diva2%3A1191313&dswid=-3643

Kerouac, J. (1999). *On the Road*. Penguin Classics.

Khurana, M. (2022). *Keeping in Touch: Situating myself amidst the incomprehensible systems of the modern world through intimate, everyday interactions* (Doctoral dissertation) [Artwork Size: 7758562 Bytes Publisher: Carnegie Mellon University]. https://doi.org/https://doi.org/10.1184/R1/21342642.v1

Kidder-Wolff, A. (2022). *The Neck Hammock : A First-Person Soma Design Project on the Experience of Rocking* (Master's thesis) [Backup Publisher: KTH, School of Electrical Engineering and Computer Science (EECS) Issue: 2022:5 Series: TRITA-EECS-EX]. KTH, School of Electrical Engineering and Computer Science (EECS). https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1635520&dswid=-4953

Kim, S., Kim, H., Lee, B., Nam, T.-J., & Lee, W. (2008). Inflatable mouse: Volume-adjustable mouse with air-pressure-sensitive input and haptic feedback. *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*, 211. https://doi.org/10.1145/1357054.1357090

Kleb, B., Wood, B., & Godby, K. (2017). Tufte-latex. Retrieved April 22, 2023, from https://www.overleaf.com/latex/templates/a-tufte-style-book-with-vdqi-title-and-contents-page/xqkjxvmrhmmp

Kolkman, F. (2015). Open Surgery. Retrieved April 19, 2023, from http://www.opensurgery.net/

Krogh, P. G., & Koskinen, I. (2020). *Drifting by Intention: Four Epistemic Traditions from within Constructive Design Research*. Springer International Publishing. https://doi.org/10.1007/978-3-030-37896-7

Kwon, H., Bae, S.-H., Kim, H., & Lee, W. (2012). Inflated roly-poly. *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*, 189–192. https://doi.org/10.1145/2148131.2148172

Let it Be! art agency. (2017). Air Slaves. Retrieved April 21, 2023, from https://we-make-money-not-art.com/air-slaves-could-we-one-day-lend-our-lung-to-filter-polluted-air/

ling QL. (2013). Mediated reality. Retrieved April 21, 2023, from https://www.designboom.com/technology/ling-ql-explores-wearable-technology-with-reality-mediators-11-20-2013/

Liu, J., Byrne, D., & Devendorf, L. (2018). Design for Collaborative Survival: An Inquiry into Human-Fungi Relationships. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3173574.3173614

Lyons, L., Slattery, B., Jimenez, P., Lopez, B., & Moher, T. (2012). Don't forget about the sweat: Effortful embodied interaction in support of learning. *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*, 77–84. https://doi.org/10.1145/2148131.2148149

Mayer-Schönberger, V. (2011). *Delete: The Virtue of Forgetting in the Digital Age* [OCLC: 1104878395]. Princeton University Press.

Merriam-Webster. (n.d.). Ruminate. Retrieved March 15, 2023, from https://www.merriam-webster.com/dictionary/rumination

Mols, I., van den Hoven, E., & Eggen, B. (2016). Informing Design for Reflection: An Overview of Current Everyday Practices. *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*, 1–10. https://doi.org/10.1145/2971485.2971494

Muthukumarana, S., Messerschmidt, M. A., Matthies, D. J., Steimle, J., Scholl, P. M., & Nanayakkara, S. (2021). ClothTiles: A Prototyping Platform to Fabricate Customized Actuators on Clothing using 3D Printing and Shape-Memory Alloys. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–12. https://doi.org/10.1145/3411764.3445613

Nadel, L., Campbell, J., & Ryan, L. (2007). Autobiographical Memory Retrieval and Hippocampal Activation as a Function of Repetition and the Passage of Time. *Neural Plasticity*, *2007*, 1–14. https://doi.org/10.1155/2007/90472

National Geographic Society. (2022). Continental Drift versus Plate Tectonics. Retrieved May 7, 2023, from https://education.nationalgeographic.org/resource/continental-drift-versus-plate-tectonics/

Neustaedter, C., & Sengers, P. (2012). Autobiographical design in HCI research: Designing and learning through use-it-yourself. *Proceedings of the Designing Interactive Systems Conference*, 514–523. https://doi.org/10.1145/2317956.2318034

O'Donohue, W. T. (Ed.). (2009). *General principles and empirically supported techniques of cognitive behavior therapy*. Wiley.

Pourjafarian, N., Withana, A., Paradiso, J. A., & Steimle, J. (2019). Multi-Touch Kit: A Do-It-Yourself Technique for Capacitive Multi-Touch Sensing Using a Commodity Microcontroller. *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, 1071–1083. https://doi.org/10.1145/3332165.3347895

Roh, Y., Heo, G., & Whang, S. E. (2021). A Survey on Data Collection for Machine Learning: A Big Data - AI Integration Perspective [Conference Name: IEEE Transactions on Knowledge and Data Engineering]. *IEEE Transactions on Knowledge and Data Engineering*, *33*(4), 1328–1347. https://doi.org/10.1109/TKDE.2019.2946162

Sadokierski, Z. (2019). Critical Journal / Contextual Portfolio: A framework for documenting and disseminating RtD as scholarly research [Artwork Size: 1110912 Bytes Publisher: figshare], 1110912 Bytes. https://doi.org/10.6084/M9.FIGSHARE.7855829.V2

Schiphorst, T. (2006). Breath, skin and clothing: Using wearable technologies as an interface into ourselves [Publisher: Routledge _eprint: https://www.tandfonline.com/doi/pdf/10.1386/padm.2.2.171_1]. *International Journal of Performance Arts and Digital Media*, *2*(2), 171–186. https://doi.org/10.1386/padm.2.2.171_1

Sebald, W. G., Bell, A., & Wood, J. (2011). *Austerlitz* (2011 Modern Library trade pbk. ed., 10th anniversary ed). Modern Library.

Sellen, A. J., & Whittaker, S. (2010). Beyond total capture: A constructive critique of lifelogging. *Communications of the ACM*, *53*(5), 70–77. https://doi.org/10.1145/1735223.1735243

Suchman, L. A. (2002). Practice-Based Design of Information Systems: Notes from the Hyperdeveloped World. *The Information Society*, *18*(2), 139–144. https://doi.org/10.1080/01972240290075066

Tharp, B. M., & Tharp, S. M. (2018). *Discursive design: Critical, speculative, and alternative things*. The MIT Press.

Theado, M. (2000). *Understanding Jack Kerouac*. University of South Carolina Press. Retrieved April 24, 2023, from https://books.google.com/books?hl=en&lr=&id=few9Pq3rmk8C&oi=fnd&pg=PP13&dq=on+the+road+kerouac+writing+style&ots=a3YyKHPafJ&sig=zA7ynihFefbv8u7keG8f7WvLdzo#v=onepage&q&f=false

Trapnell, P. D., & Campbell, J. D. (1999). Private self-consciousness and the five-factor model of personality: Distinguishing rumination from reflection. *Journal of Personality and Social Psychology*, *76*(2), 284–304. https://doi.org/10.1037/0022-3514.76.2.284

Troiano, G. M., Pedersen, E. W., & Hornbæk, K. (2015). Deformable Interfaces for Performing Music. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 377–386. https://doi.org/10.1145/2702123.2702492

Tsumura, K. (2015). FINAL HOME. Retrieved April 21, 2023, from https://www.finalhome.com/post/48825540863/%E6%A3%AE%E7%BE%8E%E8%A1%93%E9%A4%A810%E5%91%A8%E5%B9%B4%E8%A8%98%E5%BF%B5%E5%B1%95love%E5%B1%95-%E6%B4%A5%E6%9D%91%E8%80%95%E4%BD%91final-home%E5%87%BA%E5%B1%95

Viny, J., & Desjardins, A. (2020). Investigating Opaque Infrastructures: The Desktop Odometer. *Companion Publication of the 2020 ACM Designing Interactive Systems Conference*, 313–316. https://doi.org/10.1145/3393914.3395821

Weinberg, G., Orth, M., & Russo, P. (2000). The embroidered musical ball: A squeezable instrument for expressive performance.

CHI '00 Extended Abstracts on Human Factors in Computing Systems, 283–284. https://doi.org/10.1145/633292.633457

Weinreich, R. (1990). *The spontaneous poetics of Jack Kerouac: A study of the fiction* (1st pbk. ed). Paragon House.

Wong, E., & Sridharan, S. (2001). Comparison of linear prediction cepstrum coefficients and mel-frequency cepstrum coefficients for language identification. *Proceedings of 2001 International Symposium on Intelligent Multimedia, Video and Speech Processing. ISIMP 2001 (IEEE Cat. No.01EX489)*, 95–98. https://doi.org/10.1109/ISIMP.2001.925340

Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research through design as a method for interaction design research in HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 493–502. https://doi.org/10.1145/1240624.1240704

Zimmerman, J., Stolterman, E., & Forlizzi, J. (2010). An analysis and critique of *Research through Design*: Towards a formalization of a research approach. *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, 310–319. https://doi.org/10.1145/1858171.1858228

Zuckerman, O., Press, V. S., Barda, E., Megidish, B., & Erel, H. (2022). Tangible Collaboration: A Human-Centered Approach for Sharing Control With an Actuated-Interface. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3491102.3517449

Zurstraßen, N. (2020). Programming a Guitar Tuner with Python. Retrieved April 19, 2023, from https://www.chciken.com/digital/signal/processing/2020/05/13/guitar-tuner.html

Appendix A: Technical Implementation

The artifact I started with was one of the original pouches (Fig. 57) I developed with Violet Han and Di Wu for Professor Lining Yao's class.

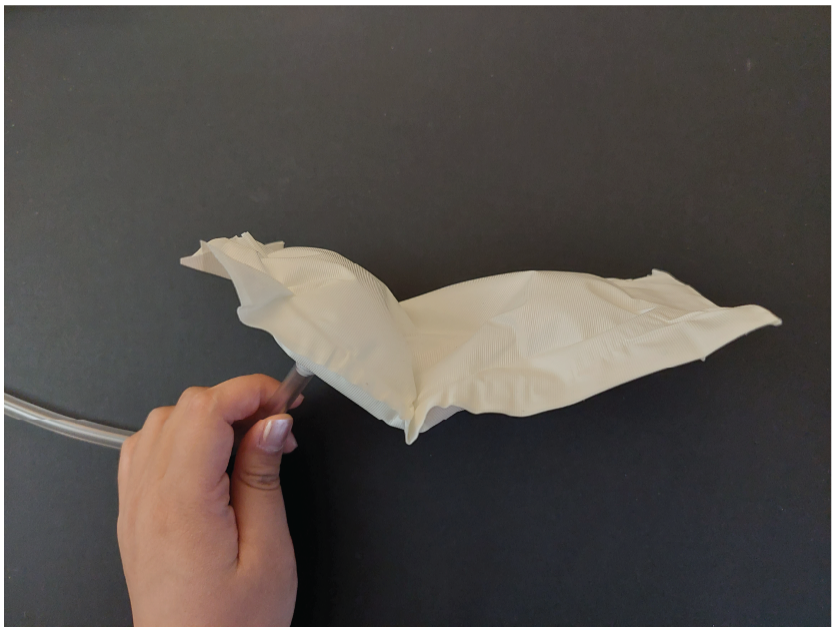


Figure 57: The original sticky pouch, developed with Violet Han and Di Wu in Professor Lining Yao's class in Spring 2022.

Arduino Capacitive Sensing and Communication

I used the Elegoo Uno R3 Board, https://www.elegoo.com/products/elegoo-uno-r3-board

Capacitive Sensing Characterizations

Since the pouches have two layers, and their shapes are determined by where the two layers make contact with each other, I resorted to

the highly customizable capacitive sensing toolkit Multi-Touch Kit (MTK) (Pourjafarian et al., 2019). This section presents the effort to study what this toolkit could do.

First, I used the code for MTK without a multiplexer to analyze how the signals from analogRead correspond to the capacitance of the system being measured.

MTK utilizes a form of capacitive sensing by measuring the voltage response of the TX electrodes, the RX electrodes, and the human finger coupled by a high-frequency oscillatory electric signal. The human body in the frequency range used can be modeled as a conductor with fixed resistance and capacitance (Bae et al., 2012). In the MTK paper, they tested different waveforms and frequencies of transmitting signals to get the maximum voltage difference in response to a touch event. I'm using their final result of a 4MHz, 25

The circuit is shown in Fig. 58.

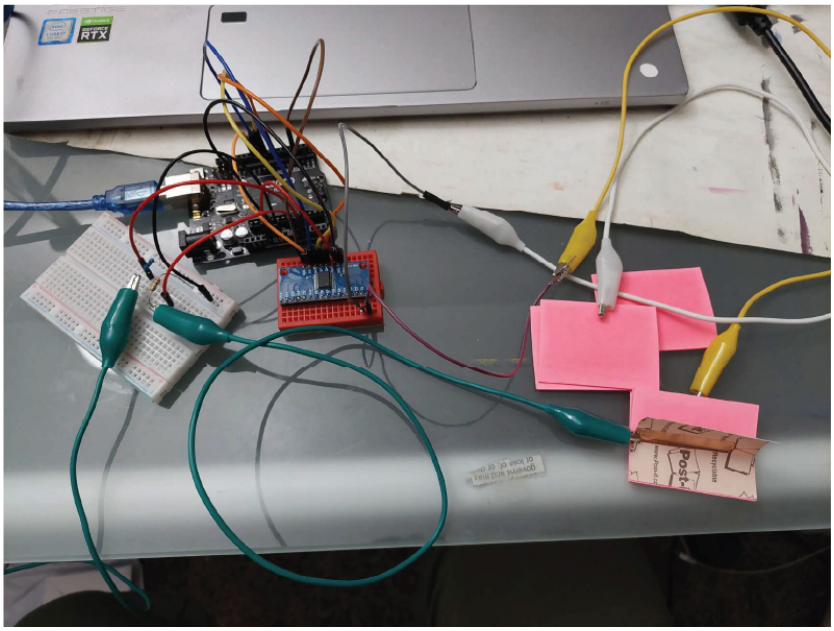


Figure 58: The circuit and the post-it notes used to characterize MTK signal.

I used copper tapes as TX and RX electrodes with varying numbers of post-it notes between them or varying overlapping areas between them.

Each RX electrode is also connected to a 100k resistor, which then connects to ground. As with the multiplexer in the provided MTK code, signals for this system come directly from the digital pin 3 of the Arduino. This receives the PWM signal from the Arduino.

The example Arduino sketch and the MTK library can be found

here: https://hci.cs.uni-saarland.de/projects/multi-touch-kit/

Based on the formular for a single parallel plate capacitor, C is proportional to A/d. The results show that no matter if I touched the electrode or not, the `analogRead` signals from the MTK code are roughly linearly proportional to the capacitance of the single copper tape capacitor plus the coupled capacitance of my body, if I was touching the electrodes.

Slider Pouch

At this stage of the development, I envisioned three kinds of pouches as in Fig. 59.

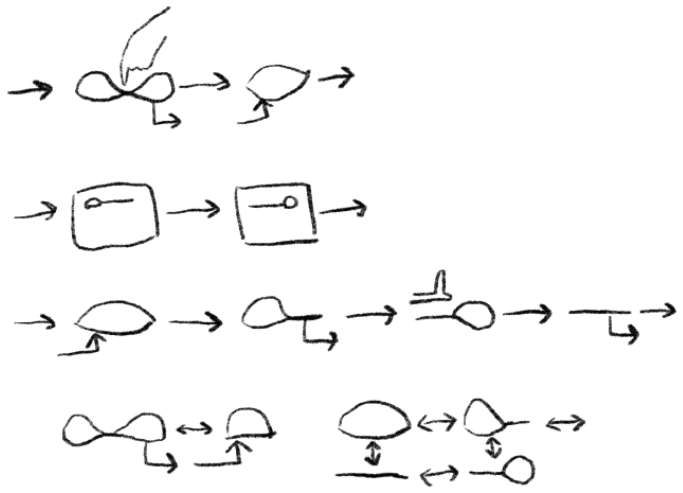


Figure 59: First row: button pouch: single point contact detection via capacitive sensing. Second row: slider pouch: linear contact detector. Last two rows: a push-pull button with 4 states.

The previous experiment could be thought of as a single-point button pouch, so I went on and implemented the slider pouch in Fig. 60.

The slider pouch itself is made out of two layers of different materials: one non-stick plastic sheet cut out of a grocery bag (on top in the image) and one fully sticky sheet of vinyl duct tape. The two sheets are held firmly in place by an embroidery hoop. On the top layer, as seen in the image, I placed a carefully cut triangle out of copper tape that tapered uniformly from one end to the other. I also put marks every 1cm on the triangle. On the bottom layer, there is a rectangular piece of copper tape, such that when the two layers are in contact with each other, the contact area is only changed by the area

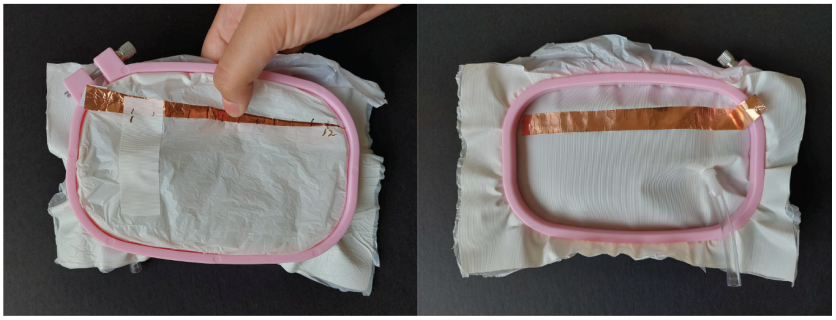


Figure 60: First row: button pouch: single point contact detection via capacitive sensing. Second row: slider pouch: linear contact detector. Last two rows: a push-pull button with 4 states.

of the triangle (explained in Fig. 61).

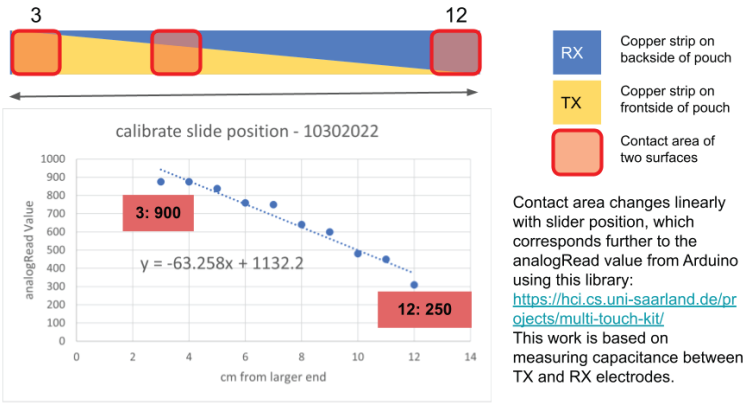


Figure 61: The characterization of the slider pouch fits the theory of capacitive sensing.

The distance from the larger end to the smaller end can be measured based on the linear decrease of the capacitance, which is proportional to the analog read values. However, this requires my finger to press into the material with the same width of the contact. This was a bit difficult to maintain, but the popping and sliding were already satisfying to me at that stage.

Slider Pouch Web Interface

I created a web page that visualizes the contact area of the slider pouch as well as the contact position on the slider with a Python[14] Flask [15] server.

This function reads the `analogRead` value from the Arduino, with the MTK code unchanged:

```
def serial_function(self):
    if self.__serial is None:
```

[14] Python 3.10.4 | packaged by conda-forge, https://www.python.org

[15] Flask 2.2.2 https://flask.palletsprojects.com/en/2.3.x/

```
            self.__serial = serial.Serial(port='COM3', baudrate=115200, timeout=.1)
        #https://stackoverflow.com/questions/1093598/
        #pyserial-how-to-read-the-last-line-sent-from-a-serial-device
        self.__serial.flushInput()
        data = self.__serial.readline() #read raw MTK println
        print(data, self.last_vals)

        data_str = data.decode("utf-8") .strip()
        list_num = re.findall(r'\d+', data_str) #tried to fix incomplete lines
        if len(list_num)>0:
            for i in list_num:
                if int(i) >= 100: #also tried to fix incomplete lines
                    val = int(i)
                    print("new", val)
                    self.last_vals.pop(0)
                    self.last_vals.append(val)
```

In the template, this function renders the contact area in the form of a ball and a slider bar position.

```
function reqListener () {
    console.log(this.responseText);
    arduinoVal = parseInt(this.responseText); //analogRead value
    ball.style.height = ((arduinoVal/1023)*130).toFixed(2) + "px";
    ball.style.width = ball.style.height;
    latest.textContent = arduinoVal;
    bar.style.width = (100-((arduinoVal-200)/850)*100).toFixed(2) + "%";
}
```

`reqListener` is called as the callback function of `var arduinoXhr = new XMLHttpRequest();`, every 100 ms. A screenshot of the webpage along with some explanations is in Fig. 62.

Multi-Touch Pouch

To capture more complex shape changes, I decided to implement the actual multi-touch sensing with concepts more similar to multiple "button" pouches. The system diagram is shown in Fig. 63.

The physical prototypes are shown in the Fig. 64.

The sensing pouch is manufactured based on the instructions provided by MTK as in Fig. 65.

The sensor I used a clear Ziploc bag with horizontal stripes of copper tape on one side and vertical stripes of copper tape on the other side. The overlapping creates capacitance between the two layers.

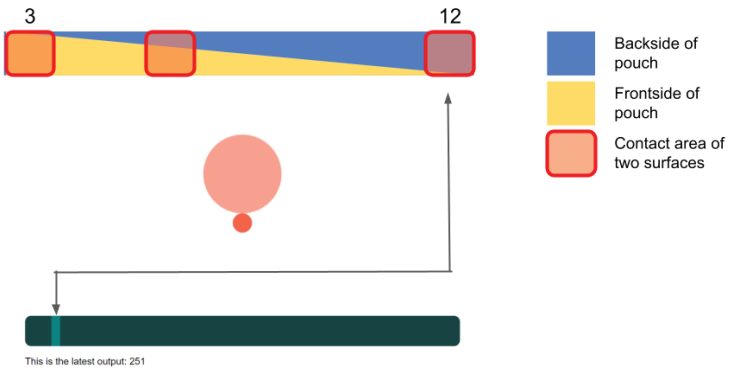


Figure 62: The web screenshot of the slider pouch with its corresponding state.

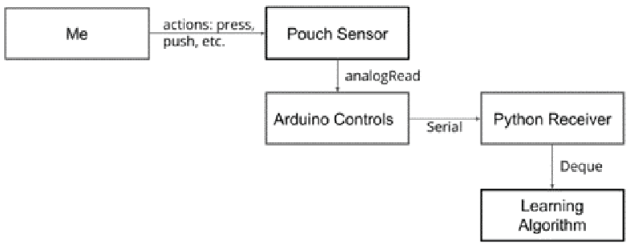


Figure 63: How a person interacts with a single pouch and it sends data to Python.

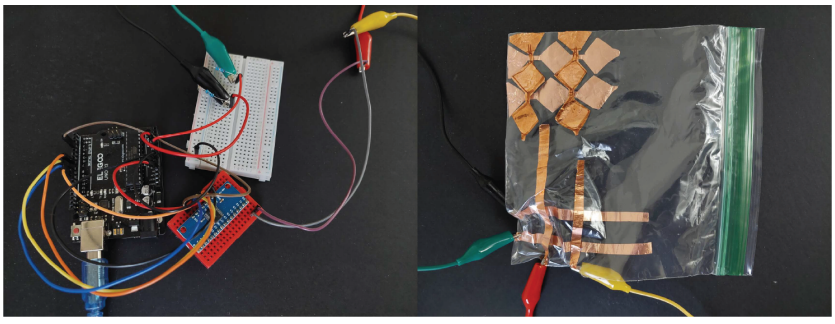


Figure 64: Left: the Arduino setup for sensing the pouch inputs. The black board is Elegoo Uno, the blue board is a multiplexer (CD74HC4607). Right: The bottom four crossing lines are the multi-touch electrodes used in this experiment. The top pattern is another option.

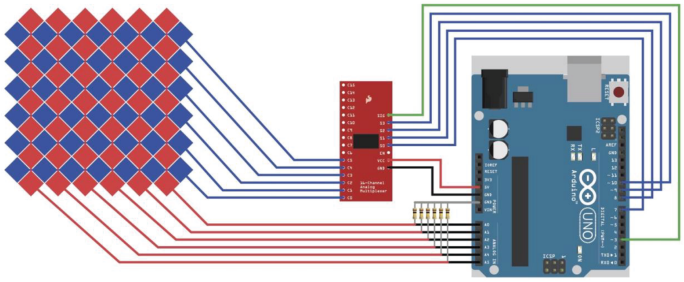


Figure 65: The circuit diagram from the Multi-Touch Kit Tutorial. Credit: Pourjafarian et al. (2019).

The circuit The vertical tapes are TX electrodes, connected to the C0 and C1 pins of the multiplexer.

The horizontal tapes are RX electrodes, connected to A0 and A1.

Each RX electrode is also connected to a 100k resistor, which then connects to ground.

On the other side of the multiplexer, VCC and GND are connected to the 5V and GND on the Arduino.

The S0, S1, S2, S3 of the multiplexer are connected to digital pins 7, 8, 9, 10 on the Arduino. Signals from these Select pins tell the multiplexer which output (C pins) to send the signal to.

The SIG (signal) pin of the multiplexer is connected to digital pin 3 of the Arduino. This receives the PWM signal from the Arduino.

Arduino Serial Print

Each intersection of the TX and RX lines gives one reading from the MTK library. To send data to my PC effectively, I added a function to the MTK library called getOneFrame, which fills an array with sensor readings instead of printing them to Serial directly.

```
void MultiTouchKit::getOneFrame(uint16_t* analog_data){
    if (Serial){
        for (int i=0; i < this->_numTx; i++){
            selectChannelOut(i);
            // Read RX
            for (int j=0; j < this->_numRx; j++){
                uint16_t val = analogRead(j);
                analog_data[i*MTK_DEFAULT_TX+j] = val;
            }
        }
    }
}
```

In the loop function of the Arduino sketch, I still print them to Serial, but all in one row that starts with "s20" and ends with "e", separated by commas. (I print 20 because without it, Serial would print characters with ASCII code corresponding to the data values instead.) For training purposes, I add a single line of "f" to mark the ending of a set of "frames", which is a set of a fixed number of readings from all TX and RX lines (called a "frame").

```
void loop() {
  mtk.getOneFrame(rawData);
  Serial.print("s");
  Serial.print(20);
  for (int i = 0; i< MTK_DEFAULT_TX*MTK_DEFAULT_RX; i++)  {
    Serial.print(",");
    Serial.print(rawData[i], DEC);
  }
  Serial.println(",e");
  counter++;
  if (counter % FRAME_LENGTH == 0) {
    Serial.println("f");
  }
}
```

Example output of a 2x2 grid of electrodes:

```
s20,70,87,78,84,e
s20,78,90,81,82,e
s20,80,89,82,92,e
s20,74,84,76,93,e
s20,77,91,80,86,e
F
```

Serial Communication to Python

I use the serial library in Python to communicate with the Arduino. `mySerial` is a subclass of `threading.Thread` such that it can be run outside the data processing pipeline. In the run function of `mySerial`, it does the following:

```
    def run_serial(self):
        #with open('C:/study/Thesis/tech/train_touch/raw_raw_data_test', 'wb') as out:
        while True:
            try:
                raw_data = self.ser.read(self.ser.inWaiting()) #read all char in buffer
                self.buffer += raw_data.decode() #decode into strings
            except UnicodeDecodeError: # catch read error and ignore it
                print('uh oh')
```

```
while '\r\n' in self.buffer: #split data line by line and store it in var
    var, self.buffer = self.buffer.split('\r\n', 1)
    if(var=="f"):
        #print("check if current data frame is complete")
        if len(self.data_buffer) == FRAME_LENGTH*MTK_DEFAULT_TX*MTK_DEFAULT_RX:
            #print("current frame complete! --------------------")
            self.deque.append(np.asarray(self.data_buffer))
            self.frame_counter+=1

            ### logging ###
            self.udp_sender.send_message("/track/arduino_frame", \
                [int(self.frame_counter)])
            if self.record_frames:
                self.frame_file.append(np.asarray(self.data_buffer))
                date= datetime.utcnow() - datetime(1970, 1, 1)
                seconds =(date.total_seconds())
                self.timestamp_file.append(np.asarray([seconds]))
            ### end logging ###
            if(self.frame_counter>self.max_frame_count):
                # terminate the serial reader after a number of lines
                print(self.line_count)
                return
        self.data_buffer = []
    else:
        self.line_count+=1
        values = var.split(',')[1:-1]
        if len(values) == MTK_DEFAULT_TX*MTK_DEFAULT_RX:
            try:
                for v in values:
                    self.data_buffer.append(int(v))
            except ValueError:
                pass
time.sleep(0.1)   #do not monopolize CPU
```

First, it reads everything in the serial buffer that Arduino has sent and tries to decode it into a string. Then, it attaches it to what is in its own buffer of strings. Once there is more than one line in the string buffer, `mySerial` splits the lines and stores them in a double-ended queue, `self.deque`. This deque is accessible to the outside from `mySerial` as clean data and has a fixed length so that only the newest data are getting trained on. Moving from passive sensing to active interaction, these attempts aimed at creating human-machine dynamics that allow the machine to support or act against the human's desire to press or release a single button pouch.

Machine Learning

Regressor

This model uses a naive simulated pouch function: the volume gain or loss at each timestep is directly proportional to the pump speed.

I first attempted to replace PID control with a machine regressor that adjusts the speed of the pumps intelligently.

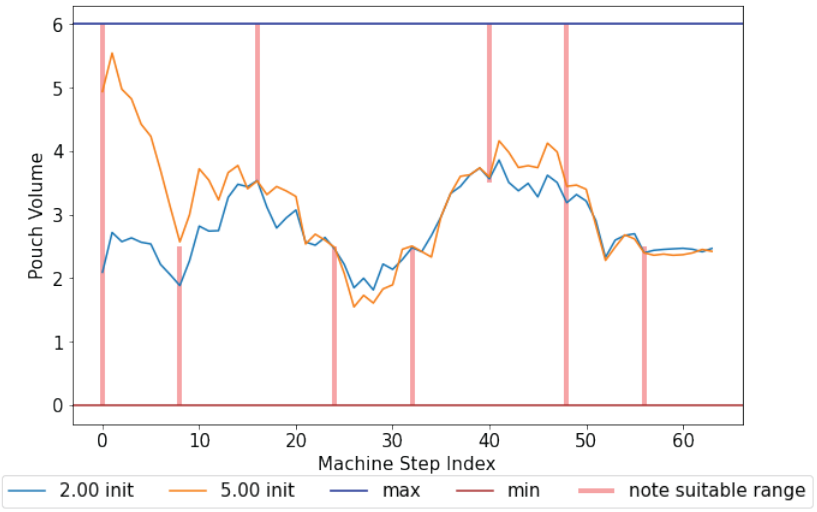


Figure 66: Output of a simulated pump with ML regression of pump speed, with l2 regularization and small gaps between acceptable ranges. The orange and blue curves are pouch volume curves as the results of the pump speed.

The model has the following loss function:

```
def VolumeRangeLossSpeed_WithNotes(y_true, y_pred, inp, step):
  initial_volume = inp[:, 0] #inp = [inital_volume, note_1, note_2, ..., note_8]
  machine_sequence = y_pred
  note_sequence = y_true
  #L2 regularization
  l2_loss = 0.03 * tf.math.reduce_sum(tf.math.square(machine_sequence))
  return l2_loss + simple_button_loss_speed(note_sequence, machine_sequence, \
    initial_volume, one_hot=True, step=step)
```

`simple_button_loss_speed` calculates how much the machine provided volume deviates from the required volume that allows the user to perform certain actions.

```
def simple_button_loss_speed(note_sequence, machine_sequence, \
        initial_volume, btn_max=6, step = 4, one_hot = False)

fixed_volume_ranges = np.array([
    [3.5, 6], #bounced back
    [0, 2.5], #can be pressed down
```

```
    [0, 6]    #can remain whatever the previous state is
], dtype = np.float32)
```

The ranges are shown as pink bars in the plot, and the loss is calculated as linearly proportional to the distance the volume is off from the required volume.

The model has the following structure:

```
__________________________________________________________________________________________
 Layer (type)                   Output Shape         Param #     Connected to
==========================================================================================
 input_27 (InputLayer)          [(None, 9)]          0           []

 dense_38 (Dense)               (None, 128)          1280        ['input_27[0][0]']

 dense_39 (Dense)               (None, 64)           8256        ['dense_38[0][0]']

 dense_40 (Dense)               (None, 64)           4160        ['dense_39[0][0]']

 input_28 (InputLayer)          [(None, 8)]          0           []

 flatten_12 (Flatten)           (None, 64)           0           ['dense_40[0][0]']

 tf.math.maximum_22 (TFOpLambda  (None, 64)          0           ['flatten_12[0][0]']

... to the loss function above
```

The regressor did a good job of meeting the required volumes, as seen in Fig. 66, but I soon realized that by the assumption that the volume difference is proportional to the pump speed, all the regressor does is to learn the slope it requires to match the next target volume. This is not realistic because the same pump speed did not mean constant air influx at all pressures, and the volume also changed at different pressures. Therefore, I went ahead and develop another pouch model.

Realistic Pouch Simulation

To model a realistic pouch, I first classified what actions the pumps and the humans could take, and what effect they have on the air in the simplest pouch with one sticky patch.

Human could: press, release

Pump could: inflate, deflate, hold, release

The volume of the pouch decreases to V_{down} when the pouch is

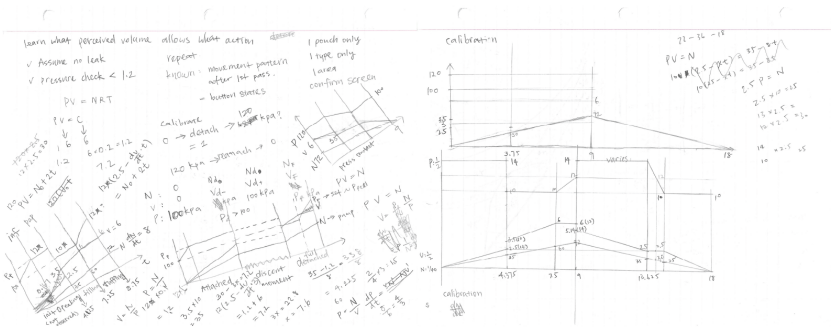


Figure 67: Earlier attempts to model the pouch dynamics with straight lines based on its states.

pressed and when the pump seals the pouch. Given that there is no air influx or release, the pressure-volume relationship is $PV = constant$. When the pump is freely pumping in air, I assumed that the pressure in the pouch is constant at P_p while the stuck area is peeling.

Based on this assumption, I drew the calibration curves as in the sketch. The calibration curve makes use of the moment the stuck area detaches from the base, detected by the capacitive sensing and caused by pump inflation, and records the characteristics of the specific pouch. This way, I would know how much to deflate to allow the user to press down fully, and how much to inflate to pop the contact area back up.

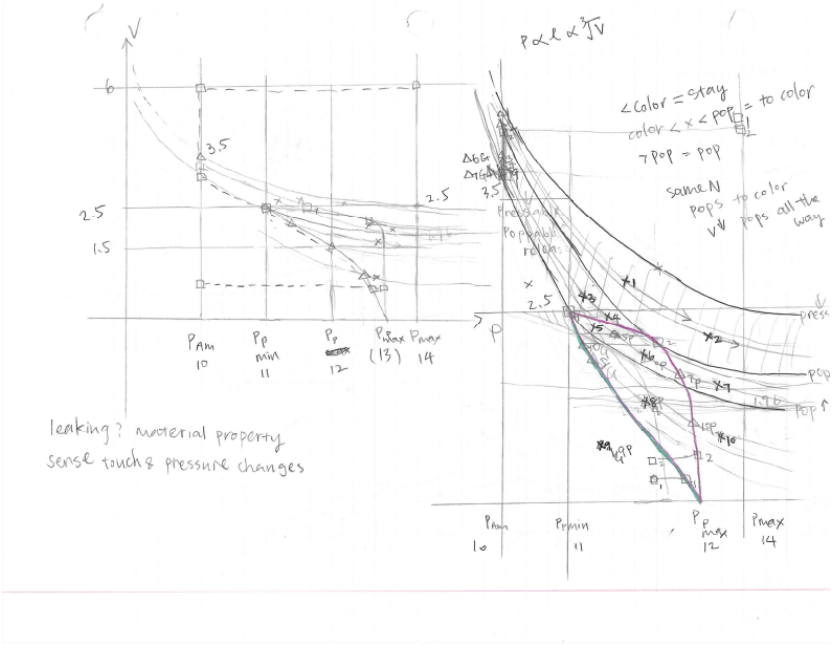


Figure 68: A more sophisticated model of the air in pouches, taking into account air inflow.

With all the manual calculations of where the curves will end, as in the breakpoints in the line plot (Fig.67), and how the inverse relationship between P and V meant that I had to draw curved lines, I drew all the actual curves in one sheet (Fig. 68). Each individual one of them corresponds to one fixed volume, and pumping air in will cause the pouch state to jump from one curve to the other.

This turned out to be intractable in the actual code, so I simplified it again for Reinforcement Learning.

RL: DQN

This model uses a more sophisticated simulated pouch function.

As I discussed before, the regressor contained a critical assumption that defeated my initial motivation to use machine learning: the pouches and the pumps are perfect. In that case, the regression result could easily be obtained as a simple required volume change / change time = pump speed (volume per second). In other words, I proved that a fully connected ml network could approximate linear segments.

This is highly unrealistic, so I created another leaky pouch function as a pump that works non-linearly depending on the pouch volume. This is not rare for real pumps. The pump:

```
def strange_factor(speed, volume):
    if speed < 0 and volume <= 0:
        return tf.cast(tf.math.exp(volume*5.),tf.float32)*float(0.2)
    if speed < 0.:
        inv_vol = 6 - volume
        return tf.math.square(tf.cast(tf.math.exp(-inv_vol/30.),tf.float32)*\
                float(1.5)*tf.cast(tf.math.cos(inv_vol/5.), tf.float32))+0.1
    return tf.math.square(tf.cast(tf.math.exp(-volume/30.),tf.float32)*\
            float(1.5)*tf.cast(tf.math.cos(volume/5.), tf.float32))+0.1
```

This "strange factor" is used as follows:

```
volume = self._prev_vol + action * strange_factor(action, self._prev_vol)
```

Where the `action` has minimum=-1., maximum=1., and represents the pump speed from full backward, to full forward. In reality, this will be controlled by two pumps and a valve. (However, this algorithm never made it out of the simulation).

In order to use a Deep Q Network (DQN) agent shipped with `tf-agents`, I discretized action space into 5 actions: `-1, -0.5, 1, 0.5, 1`.

```
env = wrappers.ActionDiscretizeWrapper( KnownVolumeEnv(button_states, volume_ranges), \
            num_actions=5)
```

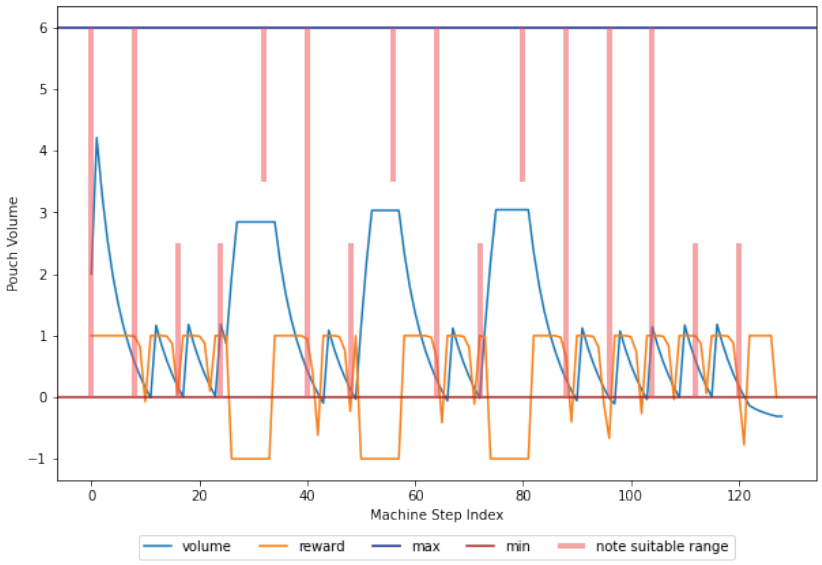


Figure 69: Output of a simulated pump with Reinforcement Learning of pump speed. The blue curve is the pouch volume curve as a result of the pump speed. The orange curve is the reward the agent receives for being within the pink target ranges.

At each timestep, the agent takes action based on the following available information:

```
prev_vol = self._state[0]
prev_speed = self._state[1]
curr_vol_min = self._state[2]
curr_vol_max = self._state[3]
next_vol_min = self._state[4]
next_vol_max = self._state[5]
```

All of these values are needed for the agent to learn the policy to act in the environment. The semi-working agent is plotted in Fig. 69. In reality, I believed that the agent would only have a rough idea of what the previous volume and speed is, and the reward (similar to the negative of the loss function for the regressor) will not be part of that agent's observation.

Due to this limitation, I realized that it is very difficult to train an RL agent to replace PID control. As a result, I investigated the possibility of using PID control, but that required some flow meters that were not readily available for such small pumps. Furthermore, that would defeat the argument of hand training a machine learning algorithm.

Pitch Detection

Pitch detection for linking the notes I sang to the pouch shapes was-based on the following tutorial. https://www.chciken.com/digital/

signal/processing/2020/05/13/guitar-tuner.html I will explain briefly what I did and what I modified.

First, I modified the function that converts frequency to a note name, instead of specifying which octave it was in. This reduces the number of classes (labels) the machine learning algorithm classifier needs to handle.

A pitch is the perceived frequency of a sound. A pure wave of 440Hz is defined as the concert pitch. This pitch is associated with the note A4.

If i is how many half steps from the concert pitch to the given pitch, we have the following formula:

$$i = 12 * \log_2(\frac{f}{f_0}), \tag{1}$$

where f_0 is 440Hz.

There are 12 half steps from the note A in an octave, so we can find the closest note from A with the following function:

```
def find_closest_note_label_only(pitch):
    i = int(np.round(np.log2(pitch/CONCERT_PITCH)*12))
    closest_note = i%12
    return closest_note
```

Now we have to detect the pitch.

I used the Harmonic Product Spectrum to help find the human-perceived pitch in the audio.

$$Y(f) = \prod_{r=1}^{R} |X(rf)|, \tag{2}$$

where $X(f)$ is the magnitude spectrum of the signal, obtained by Discrete Fourier Transform. r is the integer multiplier that aligns the harmonics of the fundamental frequency in the scaled magnitude spectrum together and cancels the noises at other locations.

[16] sounddevice 0.4.5, https://python-sounddevice.readthedocs.io/en/0.4.5

To use this method in the code, I put the pitch detection function in the callback function of Python `sounddevice`'s `InputStream`[16].

```
self.input_stream = sd.InputStream(device = 1, channels=1, callback=self.find_pitch,
                        blocksize=blocksize,
                        samplerate=samplerate)
```

The `self`, called `PitchFinder` has an instance of the `mySerial`, as `self.ser`, which fills a double-ended queue `self.ser_deque` at the rate Arduino provides, as in the previous implementation.

```
    # The sounddecive callback function
    # Provides us with new data once WINDOW_STEP samples have been fetched
    def find_pitch(self, indata, frames, time, status):
        if self.with_serial and len(self.ser_deque) < self.ser_data_length:
            # not enough Arduino to link the detected pitch to the shape signals.
            print("find pitch called with", len(self.ser_deque), "in serial")
            return
        if self.with_serial and (not self.ser.is_alive()):
            # when Arduino has quit
            self.serial_ended = True
            return
        if status:
            print(status)
        if any(indata):
            self.curr_indata = indata
            # append new samples
            self.window_samples = np.concatenate((self.window_samples,indata[:, 0]))
            # remove old samples
            self.window_samples = self.window_samples[len(indata[:, 0]):]

            signal_power = (np.linalg.norm(self.window_samples, ord=2)**2) \
                    / len(self.window_samples)
            if signal_power<POWER_THRESH:
                # the sound is too quide
                label = -1
            else:
                # avoid spectral leakage by multiplying the signal with a hann window
                hann_samples = self.window_samples * HANN_WINDOW
                magnitude_spec = abs(scipy.fftpack.fft(hann_samples)[:len(hann_samples)//2])

                # supress mains hum, set everything below 62Hz to zero
                # the noise caused by the power electricity network
                for i in range(int(62/DELTA_FREQ)):
                    magnitude_spec[i] = 0

                # calculate average energy per frequency for the octave bands
                # and suppress everything below it
                for j in range(len(OCTAVE_BANDS)-1):
                    ind_start = int(OCTAVE_BANDS[j]/DELTA_FREQ)
                    ind_end = int(OCTAVE_BANDS[j+1]/DELTA_FREQ)
                    ind_end = ind_end if len(magnitude_spec) > ind_end else len(magnitude_spec)
                    avg_energy_per_freq = (np.linalg.norm(\
                        magnitude_spec[ind_start:ind_end], ord=2)**2) \
                        / (ind_end-ind_start)
```

```
            avg_energy_per_freq = avg_energy_per_freq**0.5
            for i in range(ind_start, ind_end):
                magnitude_spec[i] = magnitude_spec[i] \
                if magnitude_spec[i] > WHITE_NOISE_THRESH*avg_energy_per_freq \
                else 0

        # interpolate magnitude spectrum
        mag_spec_ipol = np.interp(np.arange(0, len(magnitude_spec), 1/NUM_HPS), \
                np.arange(0, len(magnitude_spec)), magnitude_spec)
        #normalize it
        mag_spec_ipol = mag_spec_ipol / np.linalg.norm(mag_spec_ipol, ord=2)

        hps_spec = copy.deepcopy(mag_spec_ipol)

        # calculate the HPS
        for i in range(NUM_HPS):
            tmp_hps_spec = np.multiply(hps_spec[:int(np.ceil(len(mag_spec_ipol)\
                /(i+1)))], mag_spec_ipol[::(i+1)])
            if not any(tmp_hps_spec):
                break
            hps_spec = tmp_hps_spec

        max_ind = np.argmax(hps_spec)
        max_freq = max_ind * (SAMPLE_FREQ/WINDOW_SIZE) / NUM_HPS

        ## I drop the octave information
        label = find_closest_note_label_only(max_freq)
        label = int(label)

    # store the note in raw data queue
    self.pitch_deque.append(label/12)
    note_name = ALL_NOTES[label] if label>=0 else "-"
    if self.with_serial:
        ser_data = np.array(self.ser_deque)
        print("find pitch ser data pitch", label, note_name, \
                "max sensor reading", np.max(ser_data))
        # store the note name with the Arduino signals as a sample in the data queue
        labels = np.ones((ser_data.shape[0], 1))*label
        labaled_data = np.hstack([labels, ser_data])
        self.labeled_deque.append(labaled_data)
    else:
        # in testing mode, when there is no Arduino, simply stores the note
        self.labeled_deque.append(label)
```

ML Classification

The ML Classifier, `MyTrainer`, links the pitch data with the MTK signals and trains the network to output one of the 12 note numbers given a sequence of frames of MTK data.

For each labeled dataset, stored in `self.deque` and provided by the pitch detection module above, I separate the part with `label != -1`, meaning there is sound to train on, from`label == -1`, meaning this part needs to be predicted.

The valid labels are stored in `real_labels`, and the corresponding data is stored in `real_features`. The following snippet handles the machine forgetting. To mitigate the effect that most of the recent data in the deque are likely to be the same note and touch pattern, I maintain a dictionary `diff_labeled_sets` of data associated with each note as they come in, and the network is always trained on this dictionary, not the incoming data only.

```
for v in range(num_outputs):
    v_list = list(real_features[real_labels==v])
    print("incoming note list", v, len(v_list))
    if len(v_list) > 0:
        if v not in self.diff_labeled_sets.keys():
            print(real_features[real_labels==v].shape)
            self.diff_labeled_sets[v] = collections.deque(
                    v_list[:self.note_update_length], maxlen=self.per_label_max_length)
        else:
            self.diff_labeled_sets[v].extend(v_list[:self.note_update_length])
    else: #trying to let the machine forget earlier notes
        if v in self.diff_labeled_sets.keys() and \
                len(self.diff_labeled_sets[v])>self.note_forget_length:
            for i in range(self.note_forget_length):
                self.diff_labeled_sets[v].popleft()
```

`note_forget_length` is the amount of stored MTK data to drop at each `trian` call. This number determines how far back the machine remembers the older touch gestures, and is crucial when the user wants to assign a different note to the same gesture. If the current note the gesture is assigned to is not being forgotten, that different note will always compete with this current note when given the same gesture.

The dictionary is split into 1:4 testing and training sets to generate an error rate.

As for the network itself, I use PyTorch[17]

[17] PyTorch 1.13.1+cu116, https://pytorch.org/

```
# Define model
```

```
class NeuralNetwork(nn.Module):
    def __init__(self, num_in, num_out, num_hidden):
        super().__init__()
        #self.flatten = nn.Flatten()
        self.linear_relu_stack = nn.Sequential(
            nn.Linear(num_in, num_hidden),
            nn.ReLU(),
            nn.Linear(num_hidden, num_out),
        )

    def forward(self, x):
        #x = self.flatten(x)
        logits = self.linear_relu_stack(x)
        return logits
```

Where in this case the `num_in` is 256, `num_hidden` is 5, `num_out` is 12.

The training progress was visualized both with `matplotlib` and also a Flask-JavaScript web interface, as in Fig. 70.

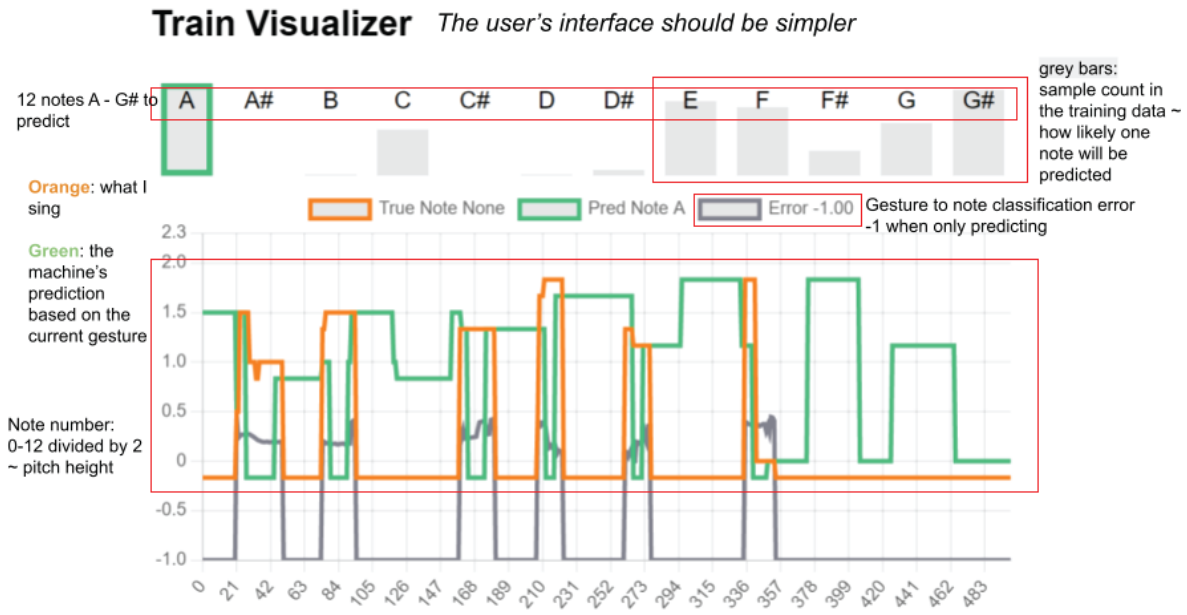


Figure 70: A web-based training visualizer with annotations.

Further Sound and Music Manipulation

Music Generation

Apart from playing the notes the machine predicts, I attempt to play some background chords that could (hopefully but not really) harmonize with the machine notes and add some texture.

I use some of the chord progressions from this website: https://www.soundgoblin.com/11-ethereal-chord-progressions-to-study/

So I wrote the roots and qualities of all the progressions and code the system to find the pattern that has the chord that contains the currently predicted note, and play that chord, and the next ones as long as possible from the same progression pattern.

For example, when the machine predicts G#, it looks for a pattern (progression) that starts with a chord that has G# in it. This starts randomly. If the next note is F#, it will first try to see if the next chord in the current progression has F#, by some "sketchy" "major +4, +9 and minor +3 +7 rule", and play it if it does, or choose another progression.

Then I use the following code for Sonic Pi[18] align the incoming notes to the beats with a soft bass beat, play the notes on the beat and discard the intermediate ones.

[18] Sonic Pi 4.3.0 https://sonic-pi.net/

The following code with comments is used to generate machine sounds, mimicking human voice, and playing background chords with Sonic Pi. A few sources are included in the comments.

```
# Welcome to Sonic Pi
# https://in-thread.sonic-pi.net/t/music-generator-integration-with-sonic-pi/6645/2
use_bpm 120
#https://raspberrypi.stackexchange.com/questions/19637/
#can-sonic-pi-play-a-continuous-midi-tone-note

# very long notes to turn on and off as requested
# synth defines what the note sounds like
machine_note = synth :dpulse, note: 30, amp: 0.3, sustain: 600
human_note = synth :growl, note: 30, amp: 0.2, sustain: 600

# collect pre-defined chords from python
live_loop :receive_chords do
  use_real_time
  key_number, quality = sync "/osc*/trigger/test_chord"
  set :chord_root, key_number
  set :chord_quality, quality
end
```

```
# collect detected human note from python
live_loop :receive_human_notes do
  use_real_time
  key_number, = sync "/osc*/trigger/test_human_note"
  set :human_note, key_number
end

# collect machine's prediction from python
live_loop :receive_machine_notes do
  use_real_time
  key_number, = sync "/osc*/trigger/test_machine_note"
  set :machine_note, key_number
end

# add a little background beat
live_loop :beat do
  #sample :drum_bass_soft, amp: 0.2
  sample :loop_3d_printer, amp: 0.02
  sleep 1
end

# play the detected note sung by a human
live_loop :play_human_sequence, sync: :beat do
  curr_human_note = get[:human_note]
  puts curr_human_note
  if curr_human_note and curr_human_note > -1
    #play curr_human_note, release: 2, attack: 0.1, amp: 1
    human_note.run
    human_note.control note: curr_human_note+12
  else
    human_note.pause
  end
  sleep 1
end

# play the predicted note
live_loop :play_machine_sequence, sync: :beat do
  use_synth :dpulse
  curr_machine_note = get[:machine_note]
  if curr_machine_note and curr_machine_note > -1
    #play curr_machine_note, release: 2, attack: 0.1, amp: 0.2
    machine_note.run
    machine_note.control note: curr_machine_note
```

```
  else
    machine_note.pause
  end
  sleep 1
end

# play the pre-defined chords received
live_loop :play_chord_sequence, sync: :beat do
  use_synth :hollow
  chord_root = get[:chord_root]
  chord_quality = get[:chord_quality]

  if chord_root and chord_root > -1
    # play the chord
    play chord(chord_root, chord_quality), release: 2, amp: 2.5
    # or play each individual notes in a sequence
    # play_pattern_timed chord(chord_root, chord_quality), 0.25, release: 2
  end
  sleep 1
End
```

Eventually I realized that because I was discarding so many notes due to predictions speed, about 90

Sound Segmentation

This module eventually replaced the `PitchFinder` module above. Named as `AudioHandler`, it identifies, segments, and saves repeated sound signals in an input stream between 1s to 3s.

At each callback from the `sounddevice InputStream`, I collect the following:

```
SAMPLE_RATE = 44100
N_MFCCS = 10
N_PER_SEG = 512
mfccs = librosa.feature.mfcc(y=snd_block[:, 0], sr=SAMPLE_RATE, \
        n_mfcc=N_MFCCS, n_fft=N_PER_SEG)
```

The callback happens every 0.2 seconds. The `mfccs` is then collected into a deque `self.segment` that stores 45 such blocks, which correspond to 9 seconds of audio data. This data is continuously updated regardless of the other functions.

From the main `plt` animation interface, `AudioHandler`'s `plot_callback` calls its `append_segment_continuous` function, which initiates the sound segmentation on whatever is available in `self.segment` at that

moment. Within `append_segment_continuous`, the `AudioHandler` first finds the assigned index in the clip library of the repeated sound identified by the algorithm, which will be explained immediately following this paragraph. It then passes the Serial data from its `MySerial` along with the clip index to `MySegTrainer` to initiate training and prediction. `MySegTrainer` shares the same structure as `MyTrainer` in the section above, but outputs a prediction array with values (probabilities) associated with all the labels.

The `assign_pattern` function segments the repeated sound and compares it with existing clips in the library to assign an index.

Within the segmentation part, I take the entire self.segment and iterate through the possible lengths a repeated sound. Since `self.segment` has 9 seconds of data, I choose to only look for clips between 1 and 3 seconds, such that there will be at least 2 repeated clips if they exist. For each fixed length, I gather all the candidates that are all the sub-clips of this length, sliding to the end of self.segment, at a small step size (5). Then I min-max normalize the candidates to be between 0 and 1, and normalize them again to a unit vector per candidate, and take the cosine similarity measure between the first segment and all the rest. This resulting line of similarity scores displays strong periodic values regardless of the actual length the current candidate has. Therefore, after proper smoothing and thresholding, I extract the median repetition starts, which are the local minima, as well as the median repetition lengths, which are the distances between local minima.

```
def get_rep_stats(self, min_seq_length, seq, candidate_step, plot_candid = False, \
            plot_products = False, plot_smooth_lines = False,):
    stats = []
    for length in range(min_seq_length, int(len(seq)/3)+1, 4):
            candidates = []
            max_candidates = np.zeros((N_MFCCS))
            min_candidates = np.zeros((N_MFCCS))
            for start in range(0, len(seq)-length+1, candidate_step):
                candidate = seq[start:start+length]
                candidate_max = np.max(candidate, axis = 0)
                candidate_min = np.min(candidate, axis = 0)
                # print(candidate[0:22])
                # norm = np.linalg.norm(candidate)
                # if(norm==0 or np.isnan(norm) ):
                if np.any(np.isnan(candidate)):
                    pass
                else:
                    # candidate /= norm
```

```
            #candidate = candidate.flatten()
            candidates.append(candidate)

        max_candidates = np.maximum(max_candidates, candidate_max)
        min_candidates = np.minimum(min_candidates, candidate_min)

    candidates = np.asarray(candidates).reshape((-1, N_MFCCS))

    # print(candidates.shape)
    candidates = (candidates-min_candidates)/(max_candidates-min_candidates)
    candidates = candidates.reshape((-1, length, N_MFCCS))
    candidates = candidates.transpose((0, 2, 1)).reshape((-1, length*N_MFCCS))
    # print(candidates.shape, candidates.max(), candidates.min())
    if plot_candid and length > min_seq_length*1.7 and \
                length < min_seq_length*1.75 :
        plt.figure()
        plt.imshow(candidates)
        print("candidates shape", candidates.shape)
        plt.show()
    if np.any(np.isnan(candidates) ):
        print("candidates nan", length)
        exit()
    #print(candidates.shape)
    candidates /= np.linalg.norm(candidates, axis = 1, keepdims=True)
    dot_products = np.matmul(candidates, candidates.T)
    # plt.figure()
    # plt.imshow(dot_products)
    dot_products = np.clip(dot_products, a_min=-1, a_max=1)
    if np.any(np.isnan(dot_products) ) or \
            np.any(dot_products>1) or \
            np.any(dot_products<-1):
        dot_1 = dot_products[dot_products>1]
        dot_n1 = dot_products[dot_products<-1]
        print("dot invalid", length, dot_1, dot_n1)
        exit()
    #angles = np.arccos(dot_products)
    if plot_products and length > min_seq_length*2.7 and length < min_seq_length*2.75 :
        plt.figure()
        plt.imshow(dot_products)
        plt.show()
    #print(angles.max(), angles.min())
    # avg_angle = np.average(angles)
    # print(avg_angle, SIMILARITY_THRESHOLD)
    # average_angles.append(avg_angle)
```

```
            # same as using first candidate across all candidates I think
            line = dot_products[0]
            # find dips
            w=np.ones(int(self.block_length/INPUT_BLOCK_TIME/candidate_step),'d')
            smoothed_line = np.convolve(w/w.sum(), line, mode='valid')
            locs = signal.argrelextrema(smoothed_line, np.less)[0]
            loc_diff = locs[1:] - np.roll(locs, 1)[1:]
            if plot_smooth_lines and length > min_seq_length*1.7 \
                    and length < min_seq_length*1.75 :
                print(locs)
                print(loc_diff)
                plt.figure()
                plt.plot(smoothed_line)
                plt.show()
            if len(loc_diff) <= 1 or np.min(smoothed_line[locs])>0.93: #cosine threshold
                #print("no repetition found for ", length)
                continue
            #print(locs)
            #print(loc_diff)
            median_diff_index = np.argsort(loc_diff)[len(loc_diff)//2]
            #print(median_diff_index)
            median_diff = loc_diff[median_diff_index]
            median_loc = locs[median_diff_index]
            #print(median_diff)
            stats.append((median_diff, median_loc, np.min(smoothed_line[locs])))
    return stats
```

To assign an clip index that the machine learning model is trained on, `assign_pattern` chooses the most popular start location and length for the collected repeated clips and compares the resulting clip with the clips in `self.clip_lib`. If one existing clip is found very similar to the testing clip, but the testing clip is longer, I choose to replace the existing clip with the testing clip. Otherwise, if there is space in the `clip_lib`, I add the testing clip as a new clip with an unused index. If the `clip_lib` is full, I replace the clip being predicted by the ML model in the nearest past with the testing clip. Whenever I change what is in the `clip_lib`, a sound file with the corresponding name is saved in a folder for mixing.

Sound Mixing

The mixing of sound clips was done in Max MSP[19]. First, AudioHandler receives the output from the ML model and performs a softmax over the raw output

[19] Max MSP 8.5.2, https://cycling74.com/products/max

```
if pred_array_last is not None:
     pred_array_last = np.power(3.5, pred_array_last)/\
            np.sum(np.power(3.5, pred_array_last))*100
    int_list = [float(x) for x in  list(pred_array_last.astype(int))]
    self.udp_sender.send_message('/trigger/train_touch_mix', int_list)
```

Then the UDP receiver in Max MSP collects these integers. The annotated max patch is in Fig. 71.

Wearable Prototype

Raincoat

The two sets of copper strips cross each other under the armpit. The bottom two wrap around the body and the top two run under the arm. An image is Fig. 72.

Hat Making

To make the pouch hat, I stuck paper pieces onto a hat that is on top of a foam head model. The paper pieces were cut into random sizes and shapes. After I assembled the pieces, I numbered them and took pictures of their relative positions. Then I took the numbered pieces apart, lay them on thermoplastic sheets, and drew their contours (Fig. 73). There were 14 pieces of paper, but I only had 8 T connectors, so I only made 8 pouches while hoping to have single-layer pieces to join them. I then cut out the thermoplastic pieces with added margins, made a small perforation, and insert the fitting for the tubing across one layer of the thermoplastic sheet, as in Fig. 74 left.

Next, I solder the edges of the pouches as in Fig. 74 right. First, I put down a sheet of brown paper. Then I sandwich the thermoplastic sheet between two layers of parchment paper. The soldering iron was set at 330 degrees Celsius, and I traced the marked edges at a speed a bit slower than drawing a firm line on a blackboard with chalk. The parchment paper would leave marks. There was a little bit of a smell of melted plastic.

Then I put the pouches over the model's head, taped them together according to the numbers as much as possible, and connected them with silicone tubes (Fig. 75, middle and right). One fitting is connected to a short segment of the tube, then connected to the side opening of the T connector, and the longer tubes run across the T bar (Fig. 75, left).

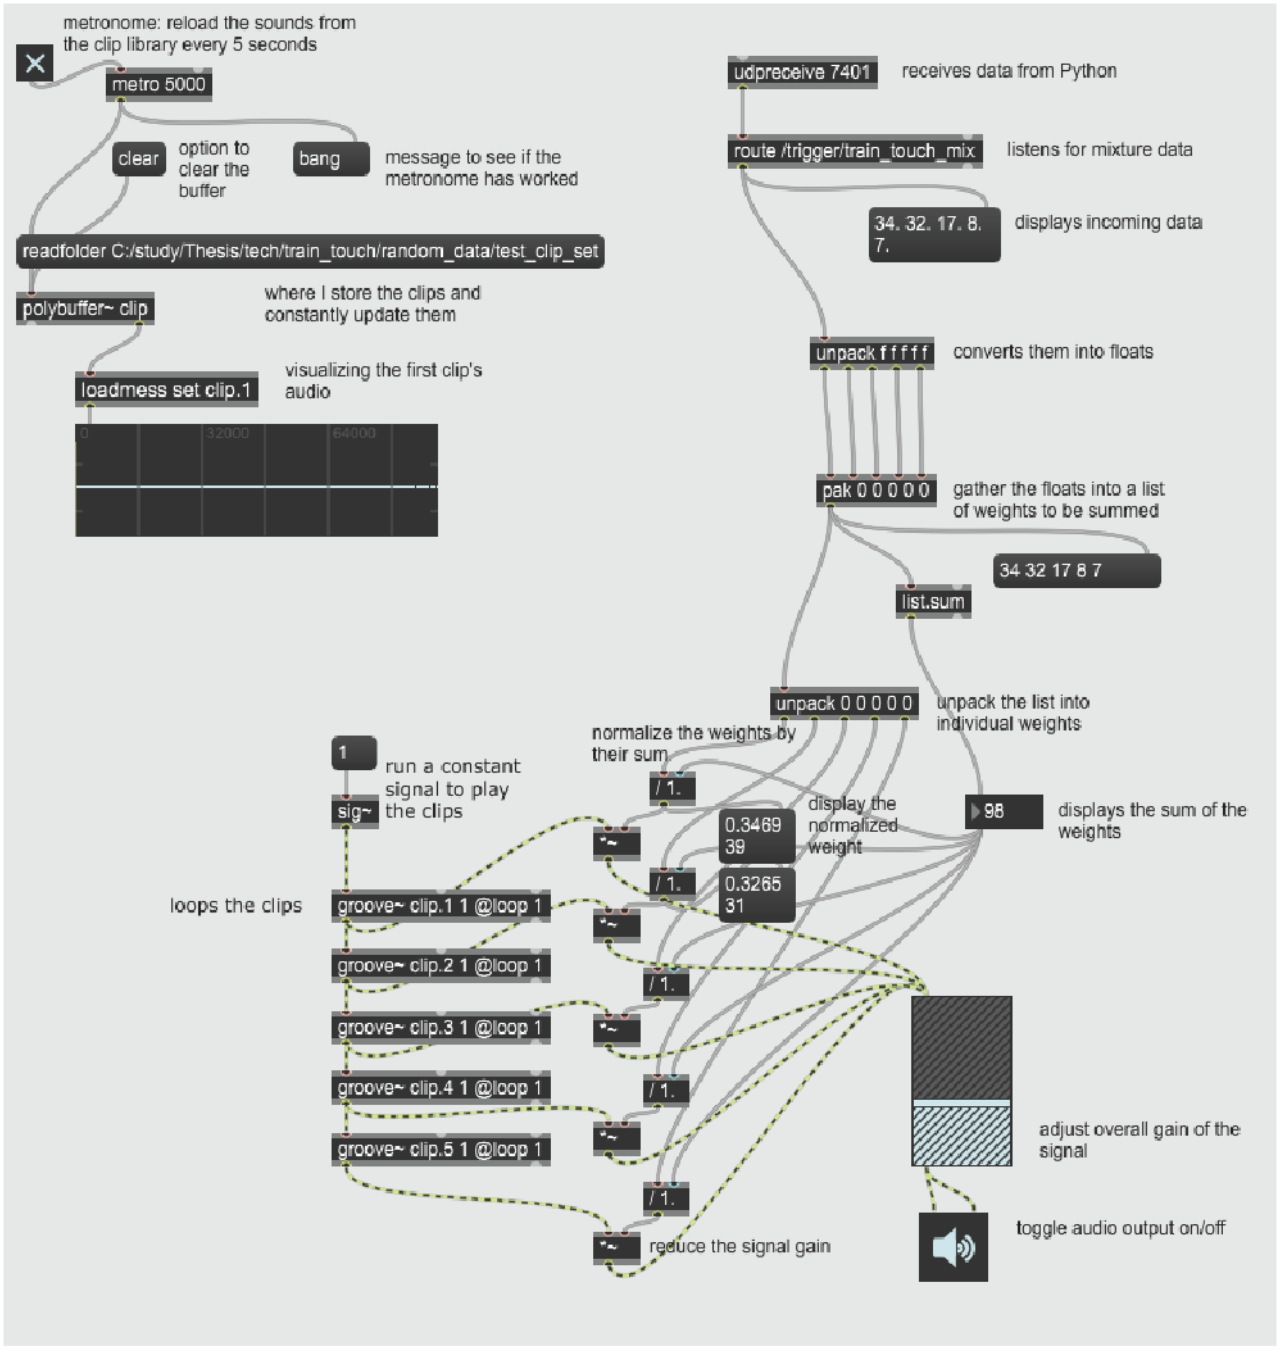


Figure 71: The annotated Max Patch that mixes the recording clips based on incoming relative weights.



Figure 72: The raincoat with crossing copper strips.



Figure 73: Left: the paper pieces were taped onto a hat. Middle and right: the pieces prepared to be made into pouches.

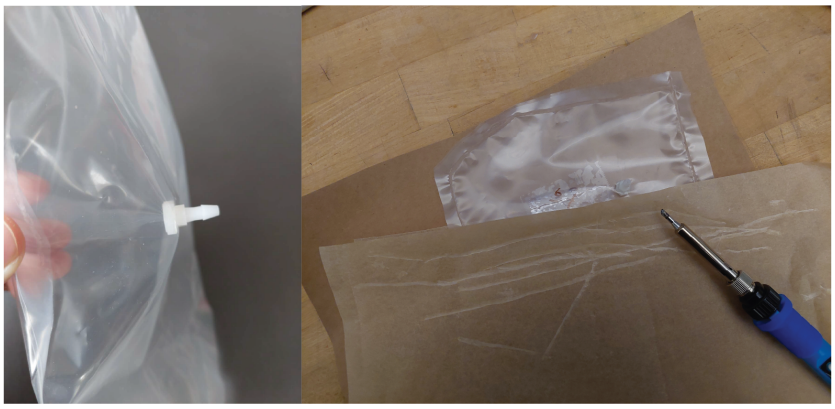


Figure 74: Left: the fitting goes through one layer of plastic. Right: the setup for sealing the pouches.



Figure 75: Left: how I used the T connector. Middle: the hat on the model. Right: the hat on its own.

Since the hat did not look the way I intended it to, I did not proceed to add the sticky patches.

Mask Making

The mask uses commercial polyethylene freezer bags for heat sealing. The overall process was identical to making the hat. The freezer bags were thinner than the plastic sheets I used for the hat, so the soldering iron was set at 280 degree Celsius. The materials and the pouch template are in Fig. 76.

The mask holds 4 pre-inflated pouches with sticky patches. I first pinched and fold the top layer of the pouch and cut out a corner from that sheet, which results in a diamond-shaped hold of a desired area. Then I put 2M red poly stucco tape (3262) over the area. The stucco tape was chosen because it peels clean from the plastic. The sticky patches then could stick to the opposite layer of the pouch and hold the distinct shapes.

The pouches were connected with tubes and T connectors in the same way the hat pouches were connected. As a result, when the user manipulates the pouches and presses into the sicky patches, the air displaced from the contact region flows into the other pouches and inflates them, enhancing the shape changes. When the user lets go of the patches, they slowly go back to equilibrium.

Adding Capacitive Sensors I put four aluminum tape strips in total as capacitive sensing strips. They are all curves made by joining the tape pieces. The insulated surface overlapped significantly, so the property of capacitive sensing still holds. In Fig. 77, I abstracted the four pouches into rectangles. The blue curves are tape curves on the outer layer of the pouches. They run over the sticky patches and serve as transmission electrodes. The green curves are tape curves on the inner side of the pouches, which cross the blue curves at where the sticky patches stick to the inner surface. This configuration is geometrically equivalent to the 4-cross layout, where two blue lines could be



Figure 76: The shape and the bag I used to make the Echo Mask.

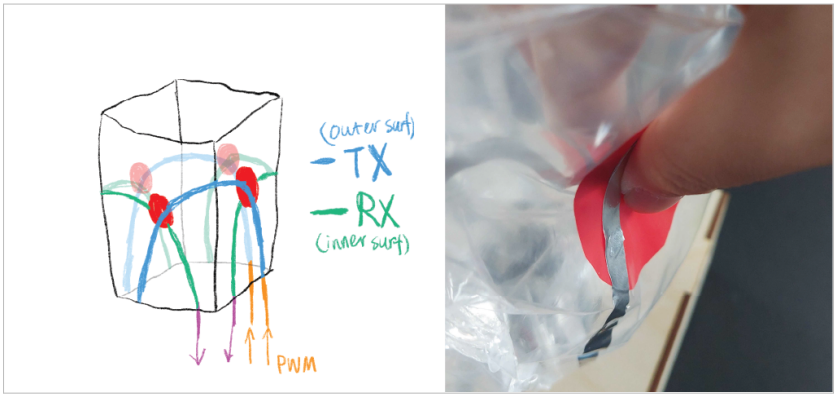


Figure 77: Left: the anatomy of the four pouches and the layout of their sensors. Right: how the sensors make contact under a pinch.

thought of as horizontal lines and the green lines are vertical crossing lines.

Box Making

Finally, I laser cut 3mm plywood pieces to hold the electronic components inside. Some of the cut patterns are in Fig. 78.

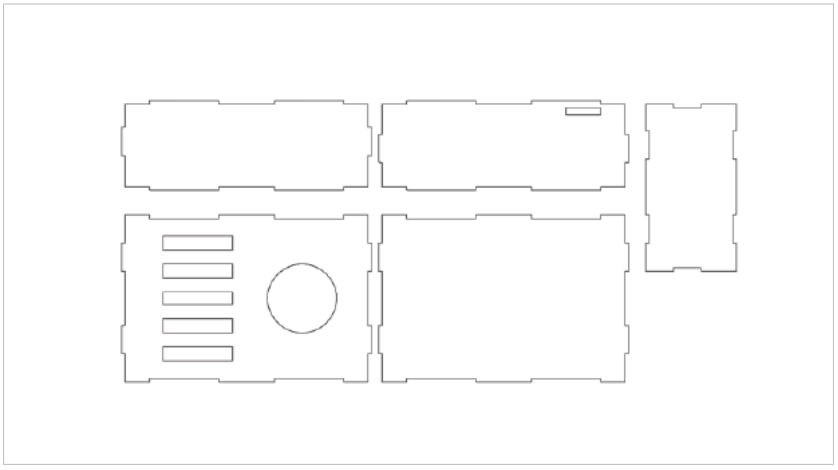


Figure 78: Some of the cut lines for the box.

The component breakdown is in Fig. 79.

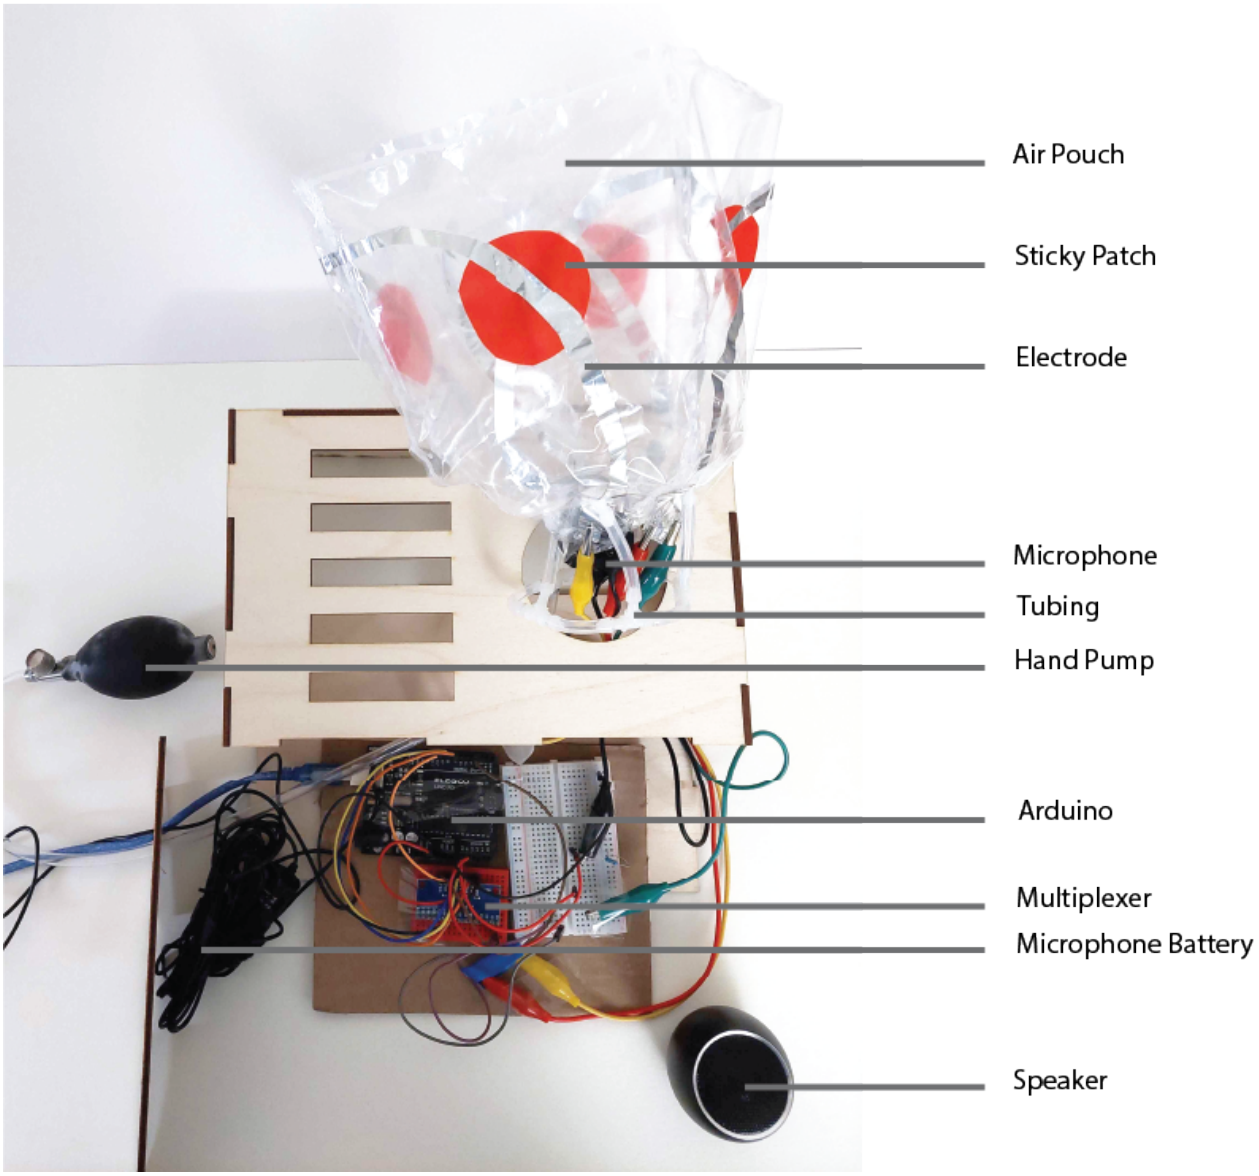


Figure 79: Annotated components of the Echo Mask.

Appendix B: A Day of a Transparent Raincoat by the Window



The Transparent Raincoat, known as TR, stands by the east facing window of a girl's apartment. To be exact, he usually leans on the windowsill, barely holding his shape. The window is right by the top of the girl's bed. The girl always keeps the blinds open, which gives

TR the first ray of sunlight every morning. Sometimes the sun shines directly at him. Sometimes the light is more scattered, softer and more ambiguous. When TR stands directly in the sunlight, his PEVA fabric distorts and reflects the scenery outside in all directions, and bits of sky often fall onto the girl's closed eyes. On those days when the sun is unsure, uncertain, indecisive on where to illuminate, TR becomes more transparent, too. You can even see through his body at the girl's desk on the other side of the window.

The girl almost always wakes up before her alarm, at 8:30 am. The first thing she sees is her phone, of course, to turn off the alarm, then she sees TR by the window. As part of her morning routine, she sits up, kicks off the covers, stands on the floor, and approaches TR. TR is not facing her. The girl picks up the hood of TR and turns him towards her. He is a delicate piece, with all sorts of wires and sensors embroidered into his fabric, so the girl's actions are gentle. "Oh my," the girl often says, as if the complexity of TR surprises her every day.

After she turns him around, the girl finds a cord coming from TR's hem and plugs TR into a machine on the desk. The machine starts to flash all sorts of lights frantically for a couple of seconds, mostly orange, but also red and blue. Then only an orange light remains lit, flickering. The girl walks into TR and puts him on.

Her hands are not touching TR. She floats freely inside him, his transparent hood covering most of her face. You can imagine that TR being a not so breathable raincoat, the girl may feel like an astronaut. Into a microphone clipped onto TR's collar, she sings a couple words to a short melody. Some examples would make this behavior much easier for you to understand. On February 3, 2023, the girl sang: I want to sleep. The melody was: so la so mi. On February 4, 2023, the girl sang: I'll cook today. The melody was: do re mi re do. On February 5, 2023, the girl sang: go buy apples. The melody was: do do mi do. These are called "phrases".

The girl repeatedly sang her words until TR is repeating after her when she pauses. TR has a hoarse voice that is strangely reminiscent of the girl's, played through the mediocre speakers on that machine. With no pauses, TR repeats the phrase. Carefully, the girl takes TR off and rests his hood on the window again.

After that, the girl goes off into her day. She brushes her teeth; gets breakfast or no breakfast; goes to work or goes to classes; buys lunch or comes home to cook; runs back to work or to school; buys dinner or cooks dinner; showers; and then she sits on her bed by TR again.

However, from time to time throughout the day, the girl comes to TR, enters him, and sings a couple words. On February 4, 2023, at 8:50 am, she got back to TR after brushing her teeth and sang: the

toothpaste is half gone. The melody was: do so mi do re do. The way she does these will seem peculiar to you.

Let's zoom into how this phrase is registered. When the girl comes back to TR after she leans him against the window, TR is still repeating the phrase. I'll cook today. do re mi re do. TR sings, somewhat mimicking the girl's voice. As carefully as always, the girl picks him up and puts him on, his hood covering most of her face. Among the many red circular pieces on TR's otherwise transparent fabric, the girl presses the one on her left shoulder into her skin.

TR is constructed with many air-filled chambers between two transparent layers. Over and under the layers are many sensing threads. The red dots on the outer layer have a sticky side facing into the chambers. When the girl presses on the red dots, they stick to the inner layer, and create a dent in their corresponding chambers.

In that particular moment, the girl sang: the toothpaste is half gone, after she released her hand from the dot and was floating freely in TR again. Hearing the girl's voice, TR came to the awareness that something had changed. At first, he was not sure, for he only knew the first phrase of the day, so he still repeated it; slowly he realized that this was a different phrase and started to register it to his memory. What he felt, the changed configuration of the red dots on his fabric, was ineffable for him. It was the girl's voice that guided him to discern his current state as distinct from the previous one, and he learned: the toothpaste is half gone. do so mi do re do.

There are about 20 dots on TR. For each phrase, the girl chooses different dots to press down. TR can remember about ten phrases like this. These ten phrases are usually more than enough for the girl for each day. On February 4, 2023, at 10:10 am, the girl pressed down the dot on her stomach and sang: gonna be late. Mi mi re do. On February 4, 2023, at 12:31 pm, the girl pressed down the dot on her left arm and sang: don't want to cook. Do ti la so. On February 4, 2023, at 3 pm, the girl pressed down the dot on the center of her chest and sang: already hungry. La ti do la so. On February 4, 2023, at 6:05 pm, the girl pressed down the dot on the right side of her waist and sang: put too much salt. Do do do do. And so on.

Sometimes the girl comes back only to listen to TR. She picks him up with care, wraps him around her, pulls down the transparent hood over her face, and listens. TR is always singing, in his somewhat muffled copy of the girl's voice, repeating what the girl has taught him. For each different set of dots that are pressed in, TR has learned the phrase associated with it. On February 4, 2023, at 4:12 pm, the girl came back to TR. He was singing: already hungry. La ti do la so. She listened for a while and detached the dot on the center

of her chest from the inner layer. TR thus landed in the previous configuration and sang: don't want to cook. Do ti la so. Sometimes TR would become unsure, as the sensors might have shifted and felt different, then he would sing a confused mix of other phrases. On that day, the girl had to repeat with the dot on her left arm pressed down this phrase: don't want to cook, do ti la so, until TR was singing as clear as he could.

When the girl is inside TR, she would be standing by the window, with her back against it. During the bright days, you can see a rim of light if you look at her backlit face and the glossy hood. When the floor lamp is on, you can see the girl has her eyes closed. She may be feeling or manipulating the dots on TR's fabric. She may be singing. She may just be standing there and listening to TR's voice through the speakers on her nearby desk.

When the girl is ready for sleep, she puts TR on one last time. One by one, she pops the dots that have been pressed down into the chamber by gently massaging the nearby area of the dots and forcing the air to invade into the attached area under the dots. Pop! It sounds. And TR goes back to his previous phrase. On February 4, 2023, they started with: put too much salt. Do do do do. Pop! already hungry. La ti do la so. Pop! don't want to cook. Do ti la so. Pop! gonna be late. Mi mi re do. Pop! the toothpaste is half gone. Do so mi do re do. Pop! I'll cook today. Do re mi re do. Then there were no more phrases. The girl usually lingers there for a while. Sometimes she plays them back all over again, rapidly pressing those dots onto her body and popping them. Sometimes she messes with TR, pressing down dots she has not touched in the day, and TR would sing an awful mix of the phrases. She would laugh lightly.

Then the girl is actually ready to sleep. She takes off TR and walks towards the machine, whose orange light has been flickering throughout the day, and unplugs TR. The light fades out. The girl leans TR on her window, making sure he is facing the outside. She lies down, pulls up the cover, and closes her eyes. TR has already forgotten her day. But she has not.

Appendix C: Interview Questions

Discussion before Sharing Document

1. Summary: what your thesis was (in the end)
2. Did you drift?
   (a) "those actions that take design away from its original brief or question and lead to a result that was not anticipated in the beginning [drifting, preface]".
   (b) Examples (location): prompted by discovery through methods, artifacts interpretations, argument not making sense, co-design process
   (c) Examples (source): internal motivation, others' feedback, literature
   (d) Can you give some examples?
   (e) What were your typical driving forces of drifts in your thesis?
3. Was that what you were used to doing? What do you think was unique to design research? If at all.
4. Will it drift more or less in autobiographical design? Do you think it makes it more legitimate, or the reverse?
5. Do you see drifts as values or drawbacks to avoid (if possible)?
6. How did you feel when facing drifts/decisions/uncertainties?
7. Did you address the drifits in your final document? If yes, how? If no, did you post-rationalize the journey, or what else?
8. Do you think it should be documented more in other forms of publications?
9. What would you like to be called in my thesis?
10. What are you feeling talking about these drifts now?

Shared Document and Questions

Thank you so much for your willingness to participate!

Here I provide 3 pages from the Journey part of my thesis. Feel free to do a rough read, and I would appreciate it if you could answer the following three questions as much as you would like.

> *[...]*
>
> *Okay, after all these, I made Sonic Pi align the incoming notes to the beats with a soft bass beat, play the notes on beat and discard the intermediate ones. Then I realized that because we are discarding so many notes, about 90% of them, the progressions did not make any difference at all (the sequence is mostly lost) so it was nearly random.*

which turned out to be a little bit like a haunted house effect. It successfully influenced my perception of the system. When I sang to my device, I could only sing gloomy and moody tunes then.

> *[1/23] Strangely, it still conveys a very strong vibe, despite the randomness and low volume, when the chord track was on, it was very hard for to deviate from 1) the note that it's playing, due to the tendency to harmonize? 2) the overall spooky, dark, weird mood, even when I want to sing some brighter melodies, like C E G major triads would be difficult. Maybe it's just me that get easily swayed into the background music?*
>
> *Why would these chords be so magical?*

I had a lot of trouble with recording the touch pouches, which were not inflated and not on-body, me singing, and the machine sounds. In what I achieved to show Professor Jesse Stiles, I was singing and tapping on a piece of dubiously looking, partly shiny piece of plastic fabric laid over my arm in a very small frame on the upper right corner of the video, and the main screen showed floating curves and bar charts as diagnostics (Fig. 27),

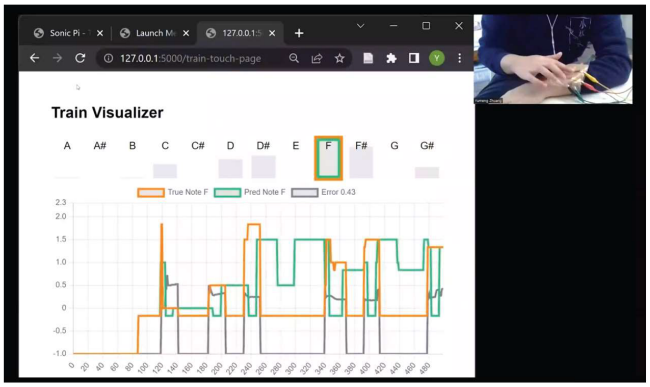


Figure 27: What I showed to Professor Jesse Stiles.

and a distant humming of synths and chords could be heard with great care. It took me about 15 minutes to explain what this mysterious Zoom recording was about, with all the concepts of hidden labor in ML, intimate musical performance, and so on. Eventually, he arrived

at the conclusion that I was trying to make a digital musical instrument with an unconventional input interface.

In reality, when I showed the recording of me interacting with the classifier pouch to Professor Jesse Stiles, he concluded that I was trying to make a musical instrument with pouches and machine learning. It did not occur to him that there was any effort involved. Instead, he interpreted the system as an attempt to bring a novel form of interaction to music-making.

This was of course very disappointing. Making a musical instrument did not touch on the speculative part, which according to me meant being something unusable. Nor did it touch upon the concept of effort, because it only took me 3 seconds to register the shape signal to the note I was singing.

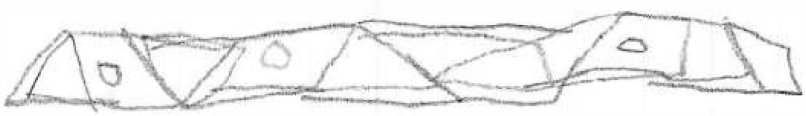
My observation through testing the system was that the ML classifier could learn the shape I provided within 3 seconds, and the difficulty was in the singing of the notes and remembering what I sang.

He did not grasp the necessity of ML anywhere, because the touch signals could possibly be analyzed with hand-crafted rules with ease. I had to agree with him that I needed to work on the argument a little more. He did introduce me to a helpful software MaxMSP, which was more powerful than Sonic Pi in generating music, and I switched to MaxMSP soon after this conversation.

Maybe because Jesse was a music professor, he did not perceive any effort in ML or singing, but when I presented this work to Malika Khurana and other students, they all perceived the difficulties in singing to the machine. To be honest, it was already hard for me to sing and hold the same note in order to register it with the shape signal, and "playing" this instrument in public was just impossible.

[1/23] For example,, I could be singing C when open, E when touch one, G when touching two. Then when I want to reinforce C, I go back to open, I would not remember what note I associated it to. I would look at that flat segment of the orange line and guess, even with a later visualization of notes with name, I am lacking the ability to see "C" and sing it. So whenever I attempted to retrain the pattern to the same note, there would be an initial slide in my voice trying to slide the orange line to where it previously was. Sometimes this was enough to overwrite the dataset buffer the ML is training on, and confused the machine to think that this pattern (touch) belongs to another note. I would then be frustrated. Especially after I visualized how fast the training dataset gets filled (and gets depleted). I had to remember the pitch! Is this a major setback in terms of usability, or is it a "feature"?

The system itself did not even work live anyway.



At a thesis mixer event, I showed people the system.

[1/23] which of course was not working, with me singing

People were self-conscious about their voices. People were not remembering what they sang and confused the ML algorithm. I should probably call this encounter of my thesis a difficult music tutor. Several people suggested that I try easier sounds,

Malika Khurana said she would not be able to sing as well and as accurately as me. She said what if some people are not good at singing, and if some people are self-conscious when singing. The singing challenge was really not something I wanted to challenge people with.

such as environmental sounds, and maybe speech, and, although I was not sure where this road would lead me,

Discussion: Hidden Labor in ML

I drifted from language generation and body gesture mapping in the previous encounter and discovered that controlling the shapes requires a lot of ML training efforts. I interpreted the required efforts as an illustration of how machine learning models often exploit human labor without people understanding the process. I proposed that my artifact could provide an explicit experience of mental and physical labor involved in training machine learning models, elicit intimate feelings towards an on-body system unique to oneself, and foster conversation about the deep implications of this technology.

In summary, with the completion of the pouch-ML pipeline, I saw the opportunity to engage the user with a critical conversation about the hidden human labor in large-scale machine learning systems. From my personal interest in raising critical questions, I was not satisfied with the concept of a musical instrument with no effort while having most effort in singing, so I drifted away from ML that was ironically too powerful to "not train". I could imagine that introducing artificial complexity in the ML model or the pouch structure could bring back the effort, but that sounded too contrived.

I used three types of writing:

The gray, right-shifted text was written in a continuous stream. I call it a "**letter**".

The left block is what I call the **main text**. It was written before the letter, in an academic style.

The typewriter-font pieces of text in the middle are excerpts from my research **journal**.

Questions:

How did you approach the text? Did you read these three styles sequentially as they appear on the page, or individually? Was it easy/hard to follow the flow?

What effect does the "letter" have on this piece of writing? Does it have a distinctive tone to you? Can you describe it?

This was one drifting moment in my development influenced by external feedback. If it were part of your thesis development process, would you document this in your thesis? If yes, how?

Any further comments?