UNIVERSITY OF COPENHAGEN
FACULTY OF SCIENCE

SHAPING MATERIAL EXPERIENCES

DESIGNING VIBROTACTILE FEEDBACK FOR ACTIVE PERCEPTION

PAUL STROHMEIER

Supervised by KASPER HORNBÆK, submitted APRIL 15TH, 2019.

This thesis has been submitted to the PhD School of the Faculty of Science, University of Copenhagen

Paul Strohmeier: *Shaping Material Experiences,* Designing Vibrotactile Feedback for Active Perception.

Supervised by Kasper Hornbæk

Assessment Committee:

Daniel Vogel *University of Waterloo*
Stefania Serafin *Aalborg University Copenhagen*
Jakob Grue Simonsen *University of Copenhagen*

PREFACE

It was a foggy night in December, 2011. I was standing outside one of Paris' train stations, keeping a lookout for two strangers I was about to meet. In the weeks leading up to this meeting, I had sent out e-mails to various tattoo artists to discuss the possibility of tattooing interactive systems on to the body. One of the few people to respond was a body modification artist, who I ended up having long e-mail discussions with. We talked online about implants and other interactive body modifications. Finally we decided to meet up in person, so, with a backpack full of electronics, I took the train to Paris to meet up with one of the world's most acclaimed and respected body modification artists.

We spent the evening sharing what we knew. I learned a lot about biocompatibility and did my best to explain the technological constraints and opportunities regarding implants that I was aware of. Towards the end of the conversation, I asked, "Can you please put a magnet in my hand?"

I had seen magnetic implants before. I had a friend who could magically collect beer bottle caps with her magnetic implants as a party trick. Magnets interested me for a different reason, though: When subjected to alternating electromagnetic fields, these magnets vibrate, allowing the implantee to feel the field.

After a short discussion, we agreed I would receive an implant the next morning. At 7:30 a.m., I found myself in a chair in a piercing studio. I had seen pictures of the procedure before and was happy to have an expert perform the implantation on me. The whole thing would have taken less than three minutes, had I not almost fainted at the site of a ⌀4 mm needle. All in all, though, the process was less unpleasant than having blood drawn.

It healed up quick enough as well. Figure 1 shows the implant in action, about a week after implantation.

INITIAL EXPERIENCES WITH AN IMPLANTED MAGNET

Initially, I was disappointed. The magnet was either deeper in my hand or weaker than the ones I had seen before. This made it impossible to pick up bottle caps (which, to be honest, I was really looking forward to doing, despite my claims that it was all in the pursuit of knowledge). However, I was also unable to feel anything interesting with it, which was even more disappointing.

At the beginning, I tried exploring what I could do with my new invisible toy. Two weeks after I got the implant, I accidentally (OK, there might have been an element of curiosity involved as well) moved

Figure 1: Party trick, using implanted magnet.

a strong magnet too close to the implant. The magnetic fields did not align, which caused the magnet inside my hand to twist. This was not a pleasant sensation (though it didn't hurt – it was more like my hand tickled from the inside). I assume this was not especially beneficial to the healing process.

After a while, I almost forgot I had the magnet, and stopped fooling around with it.

SENSING ELECTROMAGNETIC FIELDS

At the time, I had just bought a new laptop for writing my bachelor thesis on. An 11" Thinkpad. It was so small that as I rested my palms on it for typing, part of my palm extended beyond it. When the laptop's ventilation fan went on, I could feel its warm wind on my palm.

One evening, I decided to make myself warm milk. Because I was holding something in my right hand, I used my left hand to turn on the microwave. Immediately, warm wind started blowing against my palm. If I had been using an oven, this would not have surprised me, but a microwave is not supposed to generate warm wind. I put down whatever I had in my right hand and reached for the microwave again. No warm wind. With my left hand: warm wind. After a moment of confusion, it hit me: the magnet was letting me feel the microwave's electromagnetic field.

Not only was I feeling the electromagnetic field: I had also performed Pavlovian conditioning on myself to associate the vibration with warm wind, as whenever my laptop heated up, the fan would turn on. The fan made warm wind blow against my palm while simultaneously vibrating the magnet, and I ended up associating one with the other.

Once I realized that, the sensation changed slightly: it was less warm wind and more a feeling of its own, though still sort of warm-wind-like. I was able to map out where the field was. Interestingly, it was highly lopsided, expanding far into the room from one corner of the microwave, while I could hardly feel it on the other sides.

The warm wind sensation has since gone away. When there is an electromagnetic field, I feel vibration. The novelty wore off quickly. I can feel microwaves and refrigerators. Most other things I can't. Sometimes I feel a security system in a library. Once, I walked along a wall in Brooklyn and could feel a massive electromagnetic field emanating from behind the wall. I never figured out what it was coming from, though.

OPEN QUESTIONS

This process and the resulting experiences left me with a number of questions. The implanted magnet is not really useful, but it is a fascinating information channel. So how could it be made more useful? What kind of information could it provide?

Initially, I was interested in creating technology to control the haptic feedback provided by an implanted device. Could we somehow map the vibration to stimului more relevant to day-to-day life? In the years that followed, I spent a significant amount of time learning how to reduce circuit sizes and power requirements to ultimately design an interactive device which might be implanted [169]. An alternative method of inducing vibrations inside the body might be to piggy-back onto the magnet and use external hardware to actuate it. An exploration of such hardware is presented in Chapters 2 - 4.

I soon realized that the much more complex question is how to design the feedback itself. For my bachelor thesis (the one I was writing when I first felt the "warm wind") I explored using vibrotactile feedback for transmitting touch-information. I found that people using my system experienced these vibrations as symbols indicative of touch. A participant explained the experience of being remotely touched in the following way: "The thing is, we both know that she is touching me. But she maybe doesn't feel like she is really touching me, as she's just touching the robot, and I don't feel it with my body, because I just feel vibration. But just the fact the we know that she is touching me can create something..." [164]. The experience remained a symbolic one, a far cry from the direct mediation I was aiming for.

See also the corresponding video on the Mediated Touch *system:* https://youtu.be/1Bx71aYF6CA.

In the mediated touch example, as with my implant, vibration provides information. The information can be interpreted and one can assign meaning to it. However, I have since come to the conclusion that providing haptic feedback through buzzing is similar to providing visual feedback with a strobe light. It works, but it does not take advantage of the richness of the sensory channel.

In my initial experience of sensory information through the magnet, I felt warm wind. I had the closest thing I can imagine to direct information transmission. Eventually the richness of this experience faded. How could we provide haptic feedback to take advantage of the richness of our tactile perception? How can I provide feedback which is not something symbolic requiring interpretation, but something akin to access to the phenomenon itself? Working on ReFlex (Chapter 5) pointed me in a direction I thought I might find answers. I explored this direction in two studies of perception presented in Chapters 7 and 10. I further explore the idea of how to present an experience instead of a symbolic representation in Chapter 14.

ABSTRACT

Imagine running your finger over a grid. The fingertip will start vibrating as it hits each individual element. This vibration is a function of both the spacing of the grid and the speed with which you move over it. Interacting with everyday objects, if our skin is vibrated, this vibration typically occurs coupled to the speed with which we perform an action. This thesis explores haptic feedback design, based around the principle of closely coupling vibrotactile feedback to user actions.

Part I of this thesis discusses novel technologies for providing haptic feedback. I present Magnetips – a technology for tracking a magnet and providing vibrotactile feedback through the magnet over short distances – and ReFlex – a flexible smartphone with bend-input which provides vibrotactile feedback coupled to the bending motion. I present interaction-scenarios enabled by Magnetips, including interactions using an implanted magnet. Using ReFlex I present an initial exploration of input-mappings and output parameters for vibrotactile feedback from which new material experiences emerge.

Part II investigates the perception of output parameters and input mappings further. I present a magnitude estimation study which provides further insight into vibrotactile output parameters. Using qualitative in-depth interviews I present an overview of the breadth of experiences which can be created by varying the input mappings of such systems. Together the studies provide an initial systematic exploration of how to parametrically create material experiences using vibrotactile actuation.

Finally, in Part III, the findings are summarized and I explore their theoretical consequences. The results highlight the need for a theoretical framing to push the research forward. I present considerations on such a framing as a starting point for future theoretical work. Specifically I highlight that the switch in perception from vibration to material experience should receive further attention and that there is utility in considering embodied interaction at temporal scales below 100 ms; below the deliberate act. I conclude Part III by presenting an example sensory augmentation technology, which highlights how the practical results and theoretical considerations might be applied.

Forestil at din finger kører over et gitter. Fingerspidsen begynder at vibrere, når denne rammer de enkelte elementer. Vibrationen er en funktion af både afstandene i gitteret og den hastighed, hvormed du bevæger dig over den. Ved interaktion med hverdagsobjekter, hvis huden vibreres, forekommer denne vibration typisk i kobling til den hastighed, hvormed vi udfører en handling. Denne afhandling udforsker design af haptisk feedback baseret på princippet om tæt kobling af vibrotaktil feedback til brugerhandlinger.

Del I af denne afhandling diskuterer nye teknologier til at give haptisk feedback. Jeg præsenterer Magnetips – en teknologi til at spore en magnet og give vibrotaktil feedback gennem magneten over korte afstande – og ReFlex – en fleksibel smartphone med bøjnings-input, der giver vibrotaktil feedback koblet til bøjningsbevægelsen. Jeg præsenterer interaktionsscenarier muliggjort af Magnetips, herunder interaktioner ved hjælp af en implanteret magnet. Ved hjælp af ReFlex præsenterer jeg en indledende udforskning af kortlægning af input samt output parametre for vibrotaktil feedback, hvorfra nye materialeoplevelser opstår.

Del II undersøger perceptionen af outputparametre og kortlægning af input yderligere. Jeg præsenterer en størrelsesestimeringsundersøgelse, som giver yderligere indsigt i vibrotaktile outputparametre. Ved hjælp af dybdegående kvalitative interviews præsenterer jeg et overblik over bredden af oplevelser, der kan skabes ved at variere kortlægning af input af sådanne systemer. Sammen giver undersøgelserne en første systematisk udforskning af, hvordan man parametrisk skaber materielle oplevelser ved brug af vibrotaktil aktivering.

Endelig opsummerer jeg i del III resultaterne og udforsker deres teoretiske konsekvenser. Resultaterne fremhæver behovet for en teoretisk ramme for at drive forskningen videre. Jeg præsenterer overvejelser om en sådan udformning som udgangspunkt for det fremtidige teoretiske arbejde. Specifikt fremhæver jeg, at skiftet i opfattelsen fra vibrationer til materielle oplevelser bør gives yderligere opmærksomhed, og at der er ræson i at overveje kropsliggjort interaktion i størrelsesordenen 100 ms; under den bevidste handling. Jeg konkluderer del III ved at præsentere et eksempel på teknologi til sensorisk augmentation, der fremhæver hvordan de praktiske resultater og teoretiske overvejelser kan anvendes.

PUBLICATIONS

In reverse chronological order.

I refer to these papers in short as

[1] – Magnetips,
[4] – ReFlex,
[3] – Haptic Textures,
and [2] – Pulse Trains.

I present them in the above order, favoring a clear narrative over chronological accuracy.

[1] Jess McIntosh, Paul Strohmeier, Jarrod Knibbe, Sebastian Boring, and Kasper Hornbæk. "Magnetips: Combining Fingertip Track-ing and Haptic Feedback for Around-Device Interaction." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19 (2019). DOI: 10.1145/3290605.3300638. URL: https://doi.org/10.1145/3290605.3300638.

[2] Paul Strohmeier, Sebastian Boring, and Kasper Hornbæk. "From Pulse Trains to "Coloring with Vibrations" : Motion Mappings for Mid-Air Haptic Textures." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18* (2018). DOI: 10.1145/3173574.3173639. URL: https://doi.org/10.1145/3173574.3173639.

[3] Paul Strohmeier and Kasper Hornbæk. "Generating Haptic Textures with a Vibrotactile Actuator." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: ACM, 2017. ISBN: 978-1-4503-4655-9. DOI: 10.1145/3025453.3025812. URL: http://doi.acm.org/10.1145/3025453.3025812.

[4] Paul Strohmeier, Jesse Burstyn, Juan Pablo Carrascal, Vincent Levesque, and Roel Vertegaal. "ReFlex: A Flexible Smartphone with Active Haptic Feedback for Bend Input." In: *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '16. Eindhoven, Netherlands: ACM, 2016. ISBN: 978-1-4503-3582-9. DOI: 10.1145/2839462.2839494. URL: http://doi.acm.org/10.1145/2839462.2839494.

hallo rosa. du bist da

ACKNOWLEDGMENTS

There's no place like home.

I wish to thank my partner, Susanne, for her massive time investment into this thesis. Each hour I invested in writing is matched by an hour she cared for Rosa. I also wish to thank Rosa. Just for being.

I thank my sister Sophie for being who she is and doing what she does. As we are both writing up our respective theses, I'm proud that she is my sister and I am her brother. I also thank my parents, Hillary and Gerhard for their love and support. A special thanks goes to Gerhard for working hard to keep my back free, throughout the last year.

You can't make something out of nothing.

In many ways, this thesis is a continuity of work and interests that date back to long before I ever suspected that I might one day write a PhD thesis. There are a number of people who have guided or accompanied me on this way, who I wish to thank.

Looking at the completed thesis, my undergraduate time at University College Maastricht (UCM) has a surpringly strong impact on it. At UCM, Oscar van den Wijngaard and Wilfred van Dellen helped me in first orienting myself in the academic world. Jessica Messman and Ike Kamphof had a great impact on the direction my interests have taken. In fact, I can trace questions I ask and ideas I discuss in this thesis back to specific discussions with them. It is also the influence of both Jessica and Ike which makes me continuously observe my own work from a third person perspective and ask myself what it is that I am doing, when I am doing research.

I then had the opportunity to collect a wealth of applied skills building prototypes and running studies under the supervision of Roel Vertegaal at the Human Media Lab at Queen's University. This thesis has direct continuity to topics explored together with Jesse Burstyn, Antonio Gomes and Juan Pablo Carrascal.

Arriving at Copenhagen University, I was introduced to yet another research tradition by Kasper Hornbæk. Working with the Body-UI group – Aske Mottelson, Joanna Bergström, Jarrod Knibbe, Henning Pohl, Klemen Lilija and Jess McIntosh – I felt that I found my footing and voice as a researcher. Thanks also to all the other friendly faces of the HCC, especially Sebastian Boring, Stina Matthiesen, Maria Menendez Blanco, Jean-Yves Moyen, Naja Holten Mœller, Pernille Bjœrn and Xiaoyi Wang.

While my time with the Responsive Environment group at the MIT Media Lab led to output not included in this thesis, it was valuable in thinking about what draws me to the type of technologies I am interested in. Especially talking to Gershon Dublon, whose research interests match mine to an uncanny extent, was a pleasure. I thank Joe Paradiso and the entire Responsive Environment group for their hospitality.

Many hands make light work.

Aside from myself a number of people helped me create this document. Thanks to Nathan Horrowitz for wordsmithing, Ti Hoang for helping me get my figures organized, and Hamid Taramit for LaTeX support. Thanks, again, to Aske Mottelson for not only biting the bullet and defending before me – allowing me to watch, listen, and learn – but also for the Danish translation of the abstract.

Much to learn, you still have.

I'd like to extend a big *thank you* to Kasper. During my time with him, I learned to know him both as a challenging supervisor who forced me out of my comfort zone, and as a caring and thoughtful person who I appreciated having around. Kasper gave me both the space and guidance I needed to shape the direction and content of my research, while ensuring that the methodological rigour and the validity of the work met his high standards.

CONTENTS

1 INTRODUCTION

Here I present research that I conducted in the last several years, exploring haptic perception and building haptic feedback technologies. This thesis explores why the material world feels the way it feels, and how we might use such knowledge of perception in the design of interactive devices.

When you reach out and touch a large boulder, what happens that makes you think, "This part here feels smooth, while this part over there is rough and sharp"? Understanding this perceptual process might one day allow us to create a synthetic representation of that boulder which provides material experiences closely approximating those of the physical object.

I could simply reconstruct the boulder and suggest you touch the duplicate. However, just as I need not reconstruct the entire boulder for you to be able to see it (it is sufficient for me to reconstruct its visual properties in an image), it may one day be desirable to reconstruct the boulder's material properties without duplicating the entire boulder. This is currently not possible. In this thesis, I demonstrate that non-grounded vibrotactile feedback is able to approximate material properties to a much closer degree than typically assumed. I achieve this by generating vibrotactile stimuli in synchronicity with the motion of the body.

My primary interest lies not in simply replicating the physical world in a virtual world. Instead, by understanding how to haptically encode information, I want to make experiences beyond our current perceptual horizon accessible to us. Might we encode gravitational waves or the movement of the stock market in ways which are intuitively understandable to us? The studies and prototypes presented in this thesis bring us one step closer to that vision.

Not all data appear equally well suited for all sensory channels. For example, the recently increased interest in gravitational waves might be related to the new practice of presenting them as audio rather than visual graphs or figures [72].

In practice, my work oscillates between tinkering with technologies and conducting basic research on perception. To create new technology is to create new experiences. Investigating how perception works is often supported by the development of novel tools and technologies. Understanding technological constraints helps us ask relevant questions regarding perception, while understanding perception teaches us how to make better use of the opportunities that technologies provide.Thus, my technological explorations and empirical studies are closely intertwined and interdependent.

In this thesis, for sake of clarity, I separate my technological tinkering from my perception research. Part I deals with novel technologies and devices. Part II deals with studies of perception. And Part III reflects on what we have learned from all these things beyond the individual contributions of each paper.

1.1 CONTRIBUTION AND POSITIONING

Because the four manuscripts upon which this thesis is based are so different, each contains its own Related Work section. In the present section, I provide a high level overview of the context of my work, and clarify how my work is distinct from other research in the area.

Much of my work explores the use of vibrotactile feedback for rendering material experiences. Specifically, I generate vibrotactile feedback so that the frequency of discrete pulses is relative to the movement of users. I demonstrate that such vibrotactile feedback – closely coupled to human motion – can be used to render a multitude of material experiences. What distinguishes my work most from related work – such as the seminal paper by Romano and Kuchenbecker [146], and the awe-inspiring and beautiful work by Ousaid, Millet, Halliyo, Régnier, and Hayworth [134] – is that I am not aiming to replicate any specific experience and that I am not necessarily interested in realism.

Romano and Kuchenbecker [146] demonstrated that it is possible to record the vibrations of a pen moving over a textured surface and then play those vibrations back on a similar pen when it is being moved over a non-textured surface. A user writing on a glass tablet could be given the sensation they would experience if they were writing on something with a rougher texture, such as paper or canvas [36, 146]. Ousaid et al. [134] demonstrated that such high-resolution haptic systems can also be designed to operate in real time, simultaniously measuring a force and playing it back. The system presented by Ousaid et al. amplifies forces, allowing users to experience what it might be like to arm-wrestle an ant. Another impressive experience provided by that system is penetrating a drop of water with a needle. One feels the resistance of the drop's surface, and then experiences the moment that the needle enters the drop, before being sucked in by surface tension [134].

In comparison, my work is low-fi, more closely related to the parametric approach used by Kildal to simulate compliance [90]. I do not measure the properties of the real world, nor do I create hyper-realistic experiences. Instead I investigate if such experiencees might be approximated using simple heuristics. In doing so, my aim is not so much to create a specific experience, such as replicating the experience of touching a stone. Instead, I am interested in learning more about what it is that makes the stone feel the way it does. The reason for my interest in the *why* rather than the *how* is that I have the hope that understanding why we experience things the way we do will help us design novel experiences in a systematic and empirical way. By asking *why*, I aim to provide an empirical approach to exploring the importance of embodiment for interaction – in reaction to, and complementary to, existing qualitative frameworks such as those by van Dijk [43] and Klemmer [96].

My interest in the *why* is in part sparked by Peter Paul Verbeek's account of Don Ihde's relations of mediation [186]. Ihde distinguishes between "embodied mediation" and "hermeneutic mediation." According to Ihde, hermeneutic mediation occurs when we check the thermometer to determine the temperature. The thermometer provides us with a symbolic representation of the world. The thermometer might be considered an extension of the world that delivers information to us. An example of embodied mediation can be found in a dental probe. Through the dental probe, the dentist experiences features of the teeth. With the dental probe, unlike with the thermometer, the experiences are not encoded as numbers or indicators which the user interprets. One might say that the dental probe extends the body towards the world, providing the doctor with additional access.

I am fascinated by this idea of embodied mediation, and by studying the *why* of haptic perception, I provide initial steps towards an empirical approach for designing this type of embodied experience. To provide utility for HCI research, such an empirical approach needs grounding in practice. I provide this grounding by presenting general purpose implementation strategies, based on my experiences of designing haptic devices.

It should be noted that all technologies explored in this thesis are transparent, in the sense that they have the ability to augment our perception without reducing our access to the world. They are augmented reality technologies, introducing new sensations without restricting existing sensations. This is in contrast to virtual reality technologies, which encapsulate the user and block perception of the surrounding physical world. This design consideration is implicitly present in all work presented in this thesis.

1.2 SYNOPSIS

As noted above, this thesis consists of three parts, the first discussing haptic technologies, the second discussing haptic perception, and the third presenting the implications of the work.

These technologies were collaborations, drawing on the varied skillsets of each author. I am proud of these two papers not only because of their technical contributions, but also because of how beautifully my colleagues and I worked as a team.

Part I - Technologies

The first technology presented in this thesis – Magnetips (Chapter 2) – is a system for remotely tracking and actuating a magnet, either subcutaneous or attached to a fingernail. The primary contribution is to demonstrate that one might build an actuator which can be controlled at a distance in such a way that it provides an impulse at a specified time and location.

The second technology is ReFlex (Chapter 5), a flexible smartphone that can change its perceived material properties using vibrotactile feedback. The paper explores input mappings for using bending to navigate digital content. We then pair these input mappings with different haptic rendering methods and deploy them in targeting tasks. We demonstrate that certain combinations allow material experience to emerge.

These perception studies are collaborations, but they were primarily driven by me. The bulk of the work and all final decisions are mine – the good and the bad.

Part II - Perception

The devices presented in the previous section can create vibrotactile feedback with much greater precision than traditional vibration motors. In our first study of material experiences, *Generating Haptic Textures* (Chapter 7), we explore how varying the control parameters of a haptic slider affects the experience of moving it.

In a second study, we keep the vibrotactile output parameters constant, and vary the motion used to generate the vibrations. In *Pulse Trains* (Chapter 10), we present a qualitative in-depth analysis of how various motion-feedback mappings are experienced.

All papers are included as published. Minor changes are made to improve readability and to better fit the format of the thesis. Many images have been reworked and improved. Comments and additional context I wish to provide directly to these manuscripts can by found in margin notes like this one.

Part III - Implications

In part III of this thesis, I summarize the findings from the previous experiments and the implications of the technology designs. I then outline theoretical implications of this work and then present considerations and initial steps towards an empirical theory of embodied perception for HCI. I conclude by providing a simple example of a possible application, and I explain how I imagine the pragmatic findings and theoretical considerations of my work could be implemented.

1.3 OTHER MATERIAL – REACTIONS

Often, a paper that is submitted and published is disconnected from the context of the research it documents. Furthermore, the need to conform to a word count means that valuable material sometimes needs to be cut. In this thesis, I remedy these problems by adding supplementary material in the chapters directly following each of my published papers. In these added chapters, I share both the context within which I chose my research questions and some post-hoc methodological reflections and supplementary material.

After Paper 1, *Magnetips* (Chapter 2), I explain how the system might be used to interact with an implanted magnet, and provide initial data on how such interaction might be experienced.

After Paper 2, *Reflex* (Chapter 5), I report on an additional experiment which was cut from the paper after an initial round of reviews. This experiment is important, as it shaped how I approached the studies of the following two papers.

Paper 3, *Generating Haptic Textures* (Chapter 7), left me with a feeling of unease. We used an established experimental method. However, I was never completely satisfied that we chose the right method, or that we took full advantage of the method in our data analysis. In the reactions to Paper 3, I examine the analysis process from first principles and present a follow-up analysis based on this examination. I hope that this work may be useful for other HCI researchers wishing to use magnitude estimation.

For Paper 4, *Pulse Trains* (Chapter 10), I struggled to fit the richness of hours of interview data into the format of a Conference on Human Factors in Computing (CHI) paper. I therefore share some more interview excerpts which I found especially interesting.

Papers 1, 3, and 4 were conducted as part of the Body-UI research project. A tension was reconciling, on the one hand, my fascination with phenomenologists' and post-phenomenologists' discussions of the role of the body, along with my fascination with embodied interaction (as encountered at Conference on Tangible, Embedded and Embodied Interaction (TEI)), with, on the other hand, the strict empirical approach to science demanded by my supervisor, Kasper Hornbaek. In fact, not only I, but all of us in the research team were confronted by a situation in which we were researching topics relating to the body in HCI without a clear methodological or theoretical grounding. This situation led to an eclectic and delightful collection of research outputs, including on-body interaction [17], Electric Muscle Stimulation (EMS) [98], truth detection [118], eTextiles for the body [55], skin irritation as a feedback method [140], and agency [18]. This situation also led to repeated brainstorm sessions on what Body-UI is and what it should be, discussions which typically did not converge to a single idea.

Body-UI *(www.body-ui.eu) is an ERC-funded project, under the leadership of Kasper Hornbaek, at the University of Copenhagen, that aims to establish a scientific foundation for the next generations of body-based UIs.*

Within this context, I developed my own ideas and principles which I find important when considering the role of the body in HCI. This development has led to the content of Part III of this thesis, where I discuss the higher level implications of my work and how they might help us form an empirical approach to embodied interaction.

1.4 CONCURRENT WORK

The work presented in this thesis is not a collection of all of my research output to date. Rather, it presents four papers which have as an intersection a fascination for materiality, and a desire for finding better ways to take advantage of an important information channel that has thus far received relatively little attention in HCI: touch.

A separate area to which I dedicated a substantial amount of attention in recent years is that of eTextiles. I conducted a technical evaluation of a textile pressure sensor matrix designed by Maurin Donneaud. Maurin, Cedric Honnet, and I wrote a paper on this matrix which was published at Conference on New Instruments for Musical Expression (NIME) 2016 [45]. Subsequently, I collaborated with Cedric again, further exploring textile matrices. As the idea of interdigitation for optimizing such pressure sensor matrices was suggested multiple times, and was being explored in the DIY and Hackerspace communities, we conducted an empirical analysis with the support of DIKU student Victor Hakinsson. Daniel Ashbrook was instrumental in turning the final drafts of that work into a paper, and Kasper – as he does – kept me on my toes making sure that the work is methodologically sound. The paper was published at TEI 2019 [173].

See also https://3dtextiles.github.io/.

I met textile and fashion designer Rachel Freire in Shanghai, where I helped her position her work so as to make it appealing for the TEI audience [54]. Together with Sophia Brueckner, Cedric, and Jarrod, Rachel and I then conducted a successful workshop on designing eTextiles for the moving body at TEI 2018 [55].

The work on textiles included not only external collaborations. Together with colleagues from the Body-UI group, I designed a wearable sensor that can detect the dynamics of approach behavior as well as the dynamics of touch behavior. We call it zPatch and presented it as a paper at TEI 2018 [172].

I helped Jarrod Knibbe create a 56 electrode EMS stimulation sleeve which was used to explore automatic calibration EMS pose control. The resulting paper was published in the IMWUT Journal and presented at UbiComp 2018 [98]. Jarrod and I collaborated with Rachel on a second generation version of the sleeve, but it has yet to be submitted for publication.

Finally, I had the pleasure of visiting Joe Paradiso's group, Responsive Environments, at the MIT Media Lab. Together with Juliana Cherston, I built a robotic mesh-lander for grappling onto asteroids. It was deployed in a zero-gravity environment [35]. My time with Responsive Environments was important in helping me frame many of the ideas I describe in Chapter 15.

Part I

DEVICES & TECHNOLOGIES

In this section I present two projects where the technology stood in the foreground. In the first paper, *Magnetips*, I explore how a magnet can be tracked and actuated over a distance. This work provides interesting opportunities for around-device interactions, while also pointing towards novel synergies with implanted magnets.

In the second paper, *ReFlex*, I explore what new interaction modalities might emerge when our mobile devices become thin, light, and flexible. We explore bend-input mappings and how they match with different types of vibrotactile feedback. It is here where I first explore the concept of coupling vibrotactile stimuli with motion.

Building physical prototypes which are technologically novel and interesting, while robust enough to be able to run an experiment on them or demo them at a public event, requires a wide range of skill sets. For the two papers presented in this chapter, I was lucky to collaborate with two extremely talented groups of people. What these papers have in common is that the process that led to them was organic and collaborative. Both papers were only possible because of all of its authors, each of whom contributed their own unique skill set. I am proud of these two papers not only because of their technical contributions, but also because of how beautifully we worked as a team.

2

MAGNETIPS: COMBINING FINGERTIP TRACKING AND HAPTIC FEEDBACK FOR AROUND-DEVICE INTERACTION

Citation

Jess McIntosh, Paul Strohmeier, Jarrod Knibbe, Sebastian Boring,and Kasper Hornbæk. 2019. Magnetips: Combining Fingertip Tracking and Haptic Feedback for Around-Device Interaction. In: *Proceedings of the CHI Conference on Human Factors in Computing Systems.*
https://doi.org/10.1145/3290605.3300638

Abstract

Around-device interaction methods expand the available interaction space for mobile devices; however, there is currently no way to simultaneously track a user's input and provide haptic feedback at the tracked point away from the device. We present Magnetips, a simple, mobile solution for around-device tracking and mid-air haptic feedback. Magnetips combines magnetic tracking and electromagnetic feedback that works regardless of visual occlusion, through most common materials, and at a size that allows for integration with mobile devices. We demonstrate: (1) high-frequency around-device tracking and haptic feedback; (2) the accuracy and range of our tracking solution which corrects for the effects of geomagnetism, necessary for enabling mobile use; and (3) guidelines for maximising strength of haptic feedback, given a desired tracking frequency. We present technical and usability evaluations of our prototype, and demonstrate four example applications of its use.

2.1 INTRODUCTION

Research is exploring ways to extend interaction beyond the physical boundaries of our devices, through *Around-Device Interaction* (e.g., [16, 53, 97, 127, 160]). This interaction style is especially promising for mobile and wearable devices, as their small displays can limit the available space for interaction and suffer from fat-thumb occlusion [21].

Integrating these techniques with mobile devices presents a range of challenges. The tracking techniques, for example, variously suffer from occlusion (e.g,. [160]), provide only coarse or two-dimensional positions [70], or require specific lighting requirements [97]. The feedback techniques require active instrumentation of the user (e.g., [159, 185]), only work in limited directions (e.g., [198]), or have yet to be demonstrated in mobile form factors (e.g., [159, 198]).

We present Magnetips, a device that uses magnetism to enable tracking and haptic feedback for around-device interaction on mobile devices. This combination enables full 3D tracking and feedback above, around, and below the device, all within a small mobile form-factor. Magnetips is the first realisation of a combined technique for around mobile device interaction that works through clothing, when visually occluded, in any spatial direction, and with high precision. Magnetips presents the first example of magnetic tracking that accounts for the effects of the earth's magnetic field; affording a truly mobile setup with greater accuracy than previously demonstrated.

We describe (1) the implementation of Magnetips, (2) an evaluation of the tracking accuracy, both whilst the device is stationary and in motion, and (3) a psychophysics study of the user perception of the feedback, describing the signal parameters that generate the strongest haptic sensation, and the relationship between feedback strength and finger position around the device.

These evalutations show that the maximum perceivable range of feedback is 56.6 mm (at which distance tracking error is 6.38 mm when under motion), and the ideal parameters (when tracking below 83 Hz) are 12 ms signals consisting of 4 ms pulses. This allows a tracking frequency of up to 83 Hz, beyond the frequency of most displays. We show that with geomagnetism cancellation algorithm, we improve the tracking accuracy by 17.4% while the device is under motion.

2.2 MAGNETIPS

Magnetips is a device that combines tracking and haptic feedback for around-device interaction. Figure 2 shows our device, consisting of: a copper coil, four small magnetometers, an IMU, a motor driver, a power source, and a magnet on the fingernail. The magnet (10x10x1 mm disc magnet) is adhered to the users' fingernail. Magnetips can be integrated in a smart-phone case or snapped on to a smart-watch, to allow for retrofitting onto existing systems. We envision that the system may also be integrated within the device itself.

As the user moves their finger around the device, the four magnetometers collect field strength readings and estimate the position of the magnet in three dimensions. To create haptic feedback, we generate a magnetic field through the coil and alternate the polarity in order to induce vibration of the magnet. To reduce the size and weight requirements of the device, we use a coil without a magnetic core.

The high frequency of Magnetips allows feedback and tracking to be interleaved, whilst maintaining a high tracking rate (>60 Hz, the tracking rate of most devices' screens). Another benefit of a magnet-based approach is that the tracking and feedback is not challenged by

Figure 2: Magnetips consists of (a) a magnetometer array to track the magnet on the fingernail, and (b) a coil to provide haptic feedback to it. These can be used simultaneously and integrated with mobile devices (c,d) to enable interactions in the space around the watch (e).

occlusion. Magnetips works through materials (such as sleeves, pockets, and bags), electronics (through displays and internal device circuitry), and through the body. Magnetips also works in three dimensions around the device, supporting above, beside, and behind device interaction.

We include an Inertial Measurement Unit (IMU) to track the orientation of the device. This enables Magnetips to reduce the effect of geomagnetism and maintain accurate tracking in mobile scenarios. To our knowledge, this is the first work that uses the orientation of the device to compensate for errors when tracking a passive magnet in 3D.

As a result of the benefits of magnetic tracking and feedback, Magnetips affords a range of novel application scenarios. It allows tracking in all directions around the device. This enables behind device control and feedback when interacting with a smartphone, for example. This can also enable on-skin interaction and feedback beside the device. As magnetic input and output techniques are not effected by occlusion from most common materials, Magnetips affords eyes-free interaction through clothes. We discuss the technical implementation and example applications in detail later in the paper.

2.3 RELATED WORK

There have been a variety of approaches to using the space around mobile devices for interaction. Capacitive sensing has been used to extend the input abilities of smartwatches [127] and extend the sensing area of smartphones to include the area above a phone [73]. Infrared (IR) depth-sensors have been used to track regions off the device, within a line of sight of the sensors [27, 97]. Similarly, computer vision based systems have been deployed in wearable technology [160]. Alternative approaches include using the skin as an interactive area [129] or physically extending the size of mobile devices, by, for example, extending the interaction space to a watch's wristband [137, 167] or to integrate with clothing [152]. Recently, electric field sensing has enabled 3D interaction around mobile devices [101, 214]. Although there is much work on increasing the input space, there is little complementary work on providing haptic feedback at the tracked point.

2.3.1 *Mid-air Haptics*

The mobility and size constraints of wearables and smartphones present challenges when integrating a mid-air haptic feedback system. Vibration, the most common feedback modality in mobile devices, for example, requires surface contact. To be used for off-device feedback, a vibration motor needs to be worn by the user (e.g., [53]), which may encumber natural hand interactions.

Several approaches to creating mid-air haptic feedback do not require active components to be worn on the hands. Ultrahaptics uses focused ultrasound to create haptic feedback in mid-air [30]. Jets of air [174] and air vortices [158] have also been explored. While these approaches work well for larger, stationary applications, currently they lack portability, which renders them unsuitable for mobile haptic feedback. In addition, they require an un-occluded position on the surface of the device to function. In contrast, magnetism can work through materials such as glass, ceramics, plastics and non-ferromagnetic metals, as commonly used in devices.

Spelmezan et al. presented a method that provides haptic feedback using electrical arcs [159]. However, this technology has only been demonstrated to work close to the surface of the device, which enables hover interaction up to 4 mm.

Finally, in FingerFlux, Weiss et al. showed that by attaching a permanent magnet to the fingertip, attracting and repelling forces can be felt via electromagnets [194]. By using an array of electromagnets, positioned beneath a surface, Weiss et al. were able to guide a users' finger during screen-based interaction. We build on this work and tailor it to a mobile solution. We replace the array of small electromagnets with a single large coil. This design allows us to reduce the size, weight and power requirements of the FingerFlux approach. Unlike FingerFlux, we can not discernibly 'push' and 'pull' the users' finger, instead we use an em-pulse-burst to create an off-device vibration in the worn magnet.

A benefit of using a passive magnet for mid-air haptic feedback is that we can simultaneously use magnetic approaches to tracking the users' finger around the device. This tracking benefits from the same features as the haptic system: tracking in three dimensions, through materials. In the next section, we present the existing literature on magnetic tracking.

2.3.2 *Magnetic Tracking*

There is much work in the area of magnetic sensing for magnet position estimation. In Finexus [32], for example, users wore electromagnets on their fingertips to enable accurate hand tracking. We envision a system that relies on only passive instrumentation of the user, however, affording a simple, less obtrusive, untethered setup. Therefore, we will focus on passive magnet tracking.

Research has explored the use of both individual magnetometers, for a lightweight yet spatially-restricted tracking, and multiple magnetometers, for full 3D tracking.

As an example of single magnetometer use, Ketabdar et al. explored gesture tracking around the device while the user wears/holds a permanent magnet [89]. They used the euclidean norm of the magnetic field strength in order to determine the magnets' distance from the device. While this did not enable precise location tracking, it did support accurate gesture recognition proximal to the device. In similar work, Ashbrook et al. [7] devised a ring with an embedded magnet, where 1D rotations and translations of the ring could be used as an additional input modality. Radial movements around a device can also be sensed by a single magnetometer as Harrison and Hudson demonstrated in Abracadabra [70].

In a variation on the device-mounted magnetometer approach, Han et al. [69] instrumented the users' wrist with two magnetometers, in order to support 2D mid-air handwriting. In uTrack, Chen et al. [33] demonstrated 3D tracking of magnets with as few as two magnetometers.

One of the challenges with using magnetometers to track magnets is the effect of the Earth's magnetic field, which needs to be taken into account during calibration, yet changes with any change in sensor orientation [32]. This can result is significant declines in tracking accuracy as a magnet moves just a few centimeters away from the sensors [32]. There have been a number of approaches to dealing with this challenge, from restricting input to limited dimensions [69], using large magnets in close proximity to the sensors [33], to using advanced signal filtering with electromagnets [32].

We present another novel solution to this approach, by using an IMU (inertial measurement unit) in order to track changes in orientation, which can then be used to cancel the effects of geomagnetism. This allows us to maintain accurate tracking without any limit to interaction, without a strong magnet, and without the use of electromagnets.

2.4 IMPLEMENTATION

Magnetips consists of two elements, a tracking system and a haptics system. We will describe these in the following sections.

2.4.1 *Hardware*

TRACKING We use four three-axis magnetometers (LIS3DML), placed in a rectangular arrangement. Figure 3 shows the arrangement of the sensors, and the exact placements of the sensor ICs within the board (32.5×26.7 mm rectangle, small enough to fit within the dimensions of an Apple watch series 3, as depicted in Figure 2). Additionally, there is also an IMU (LSM6DS33) to track the orientation of the sensor board more accurately. The sensors are able to sample at a frequency of 1 kHz. This high frequency allows us to multiplex the haptics with the tracking, and is one of the main reasons for choosing this sensor.

HAPTICS We provide haptic feedback by creating electromagnetic pulse-bursts with a hand-wound coil. As we intended to use this technology with mobile devices, using the coil without a magnetic core is important to reduce the weight and size of the device. The coil is 45 mm

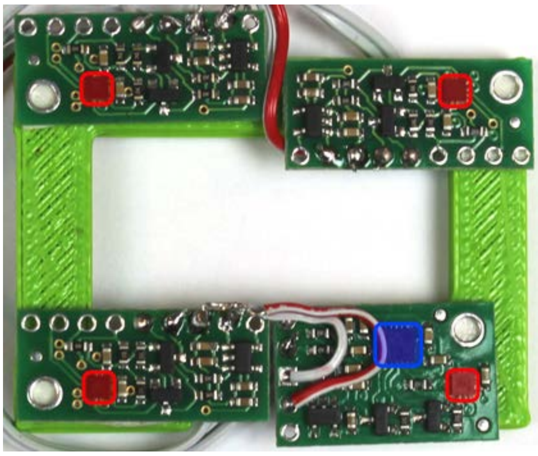

Figure 3: The array of sensors consisting of 4 magnetometers (red) and an accelerometer and gyroscope (blue).

by 55 mm and has 38 windings of 0.5 mm enamelled copper wire. The coil is driven with a Pololu VNH5019 Motor Driver Carrier. Each burst consists of a single pulse or a series of pulses with varying frequencies.

CONTROLLER Haptic feedback and sensors measurements are controlled by a Teensy 3.6. The measurements are sent to a desktop PC via serial communication, ready to be processed by a host application.

2.4.2 *Software*

TRACKING ALGORITHM We use the same algorithm used by Chen et al. in Finexus [32] to estimate the position of the magnet in 3D. This algorithm uses the magnetic field strength that each magnetometer is subject to, in order to estimate the distance between the magnet and each magnetometer by using the inverse cubic relationship of field strength to distance.

GEOMAGNETISM CANCELLATION The earth's magnetic field produces a bias in the magnetometers readings. The field consists of an inclination (the deviation between true north and magnetic north), a declination (the angle between the magnetic field lines and the earth's surface) and an intensity. These three components can be represented as transformation matrix M, which describes the rotation between a true north coordinate system and the earth's magnetic coordinate system. The field is relatively constant throughout a small geographical region, say a city, and can be calculated based on latitude and longitude[1]. This effect is used to create compasses, magnetometers are indeed digital compasses.

1 Magnetic Field Calculator: https://www.ngdc.noaa.gov/geomag-web/.

Geomagnetism, however, is a hindrance to the tracking of magnetic objects as it skews the readings of magnetometers. During interaction, the wrist of the user will move (even if slightly), thus altering the readings of the magnetometers. Simultaneously, the magnet's effect on the sensors will also be present. It then becomes challenging to know how much signal change can be attributed to movement of the magnet or movement of the device.

Our algorithm is one of the ways to overcome this challenge, which we will now describe. Firstly, the sensor's orientation (in world coordinates) has to be known. Yet, a magnetometer by itself cannot estimate the orientation of a device with 3 DoF (degrees of freedom) accurately. In the same way, an accelerometer cannot track the orientation with 3 DoF without drifting around the axis of gravity. A gyroscope can also be fused with the existing data, in order to give a more responsive estimation. Each of these sensors has 3 axes, for a total of 9 axes. A system that performs such an orientation estimation is called an AHRS (attitude and heading reference system). We implement Madgwick's AHRS to get accurate estimations of the orientation in 3 DoF.

The key to our geomagnetism cancellation algorithm is to reverse the AHRS process in order to estimate magnetometer data from the orientation data. That is, given an absolute orientation of the device, we can calculate the expected magnetometer readings (using the local earth's magnetic field components, represented as transformation M and intensity I_M) for each axis of the magnetometers. We do this as follows for each magnetometer: first, we calculate the representation of a magnetometer's axes (as vectors) in the earth's magnetic coordinate system using the earth's magnetic field transform M, the intensity I_M and the sensor orientation S_R. Second, we project these transformed vectors onto the earth's magnetic field vector (the x-axis of M). The calculations are:

$$\begin{aligned} I_{S_x} &= M^{-1} \cdot S_R \cdot (1,0,0)^T \cdot (I_M,0,0)^T \\ I_{S_y} &= M^{-1} \cdot S_R \cdot (0,1,0)^T \cdot (I_M,0,0)^T \\ I_{S_z} &= M^{-1} \cdot S_R \cdot (0,0,1)^T \cdot (I_M,0,0)^T \end{aligned}$$

Resulting values fall into $[-I_M, +I_M]$. For example, if the x-axis of a magnetometer is properly aligned with the earth's magnetic field, it will have the full intensity, while both the y-axis and z-axis will read 0. The actual magnetometer readings are then corrected using the expected readings $I_{S_{x,y,z}}$. In an ideal case (i.e., no magnetic object being present near the sensor), the corrected readings would be 0 on all axes.

However, in the presence of a magnet, 9-axis AHRS cannot be used due to the magnet's influence on the magnetometer. Fortunately, 6-axis AHRS can be used which uses only the accelerometer and gyro [110], which is sufficient for short interaction times. Over time, however, there will be a slight drift due to the absence of magnetometer data. This minor accumulated drift can be corrected for after the magnet is taken

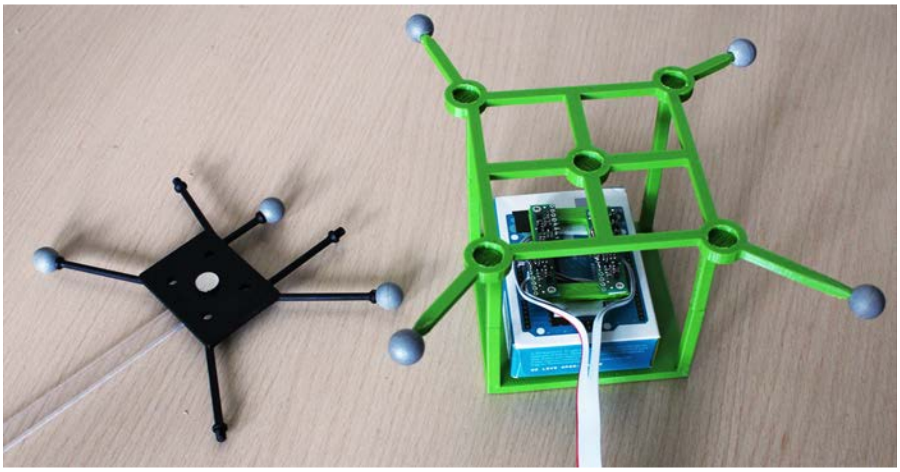

Figure 4: Setup of how the magnet and sensor board was tracked using OptiTrack.

away when the interaction is complete. One also needs to know when to switch from 9-axis to 6-axis AHRS. We implement this detection with a simple threshold on the difference between the magnitude of the earths magnetic field vector and the current measurements. The AHRS is likely to have been given some biased magnetometer data shortly before reaching this threshold. To avoid this issue, we store the 9 axes of IMU data into a buffer and the system re-calculates the 6-axis AHRS from the past second of data.

2.5 TRACKING EVALUATION

Next, we evaluate the accuracy of our tracking solution. Importantly, we demonstrate the effectiveness of our novel geomagnetism cancellation, showing an improvement of 17.4% tracking accuracy when the device is under motion. We also present results to show the interference of the magnetic field produced by our haptics system, and show that the minimum delay necessary after producing a haptic signal is 2 ms.

2.5.1 *Tracking Accuracy*

STATIC SENSORS We tested the tracking accuracy by first keeping the sensor board static and moving the magnet in a 3D volume, centred around the sensors. For this condition, we calibrated the geomagnetism before measuring the data. As the sensor base is static, our geomagnetism cancellation algorithm does not provide any further benefits. In each of these tests, Magnetips was placed on a table. The magnet was attached to retro-reflective markers for optical tracking. We moved the magnet-marker in an 80 × 80 × 80 mm volume and tracked its position with both Magnetips and OptiTrack, where OptiTrack provided the ground truth for this data (Figure 4). Our results are shown in

			Length of cube (mm)		
GC	Sensor base	Magnet	40	60	80
N/A	Static	Moving	3.001	5.055	6.686
Off	Moving	Static*	7.790		
On	**Moving**	**Static***	**5.514**		
Off	Moving	Moving	4.408	7.502	9.095
On	**Moving**	**Moving**	**4.354**	**6.379**	**7.510**

Table 1: The average errors accumulated over the duration of data sampling in mm. Geomagnetism cancellation abbreviated to GC. **Relative to the sensor base.*

table 1. We present these results with varying volume sizes. The error is calculated by the euclidean distance between the OptiTrack position and the Magnetips position (mm) for each frame of the data. The error for each frame is then averaged. The magnetometer data sampling rate was set to 1 kHz, with ±8 gauss sensitivity. The results show that there is an approximately linear increase in error from the 40 mm to the 80 mm cube lengths.

MOVING SENSORS The previous results, however, show the upper bound of tracking accuracy where the sensing base is not moving and therefore the geomagnetism is constant. Therefore, we conducted two further tests. Firstly, we fixed the magnet at a known position away from the sensor board, using the green cage shown in Figure 4 (in the middle indentation). We then recorded the data whilst moving and tilting the sensor board and magnet. In this test, the previously used algorithm performs poorly and results in an average error of 7.790 mm. After using our geomagnetism cancellation technique based on the AHRS, our findings show an average error of 5.514 mm (see table 1 for comparisons between geomagnetism cancellation on and off). This is an accuracy improvement of 2.276 mm, for a magnet position fixed at 5cm above the sensor base.

Then finally, to simulate a more realistic scenario, we moved the base and the magnet simultaneously. Again, with our geomagnetism cancellation we were able to reduce the error in every test we conducted. The mitigation technique seems to have a greater effect with a larger volume, probably due to the fact that the earths magnetic field has a stronger effect relative to the effect of the magnet when further away. In the largest volume, we see a decrease in the error of the tracking by 1.585 mm.

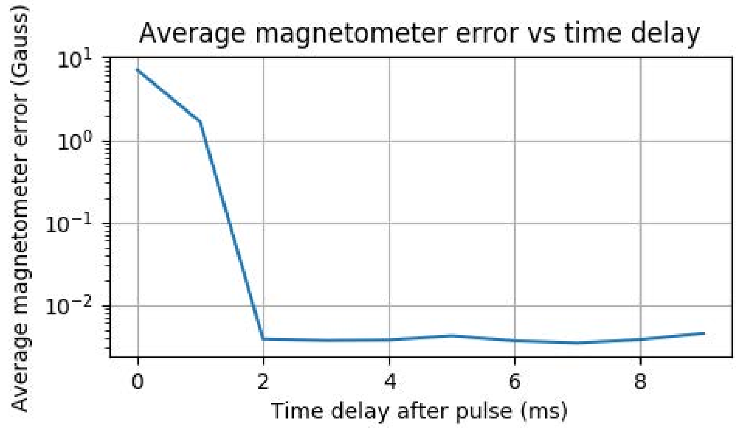

Figure 5: Plot of the average magnetometer errors vs the time since a haptic pulse had ended.

2.5.2 *Multiplexing Tracking With Haptics*

The haptics system we use inevitably interferes with the magnetometer readings without careful timing. We investigated whether there is a lasting effect from the coil after the controller turns off power to it, or whether the latency of the reading is an issue. Indeed, we found that if we begin to sample data just afterwards, there is some error in the readings as shown in Figure 5. The error in this graph represents the average error from the earths magnetic field on a logarithmic scale. The interference stops after a 2 ms delay, but it is very large and unpredictable before this. It is also necessary to leave at least a 1 ms delay after reading, before sending a haptic signal. Due to the fact that we can only use the magnetometer in continuous sampling mode at 1 kHz, this means we have to leave a 2 ms delay after reading. In total, this means that there must be a 4 ms delay added onto the sampling period. This limits the sampling frequency of the tracking if haptic pulses are to be used often during the interactions. We discuss the implications of this in the next section.

2.6 HAPTICS EVALUATION

We conducted a series of evaluations to demonstrate haptic feedback properties. Tracking accuracy can simply be measured, but the experience of haptic feedback is subjective and not as easy to quantify. Therefore, we spend more time on the feedback evaluation than on the tracking, though both parts are equally important for Magnetips.

Although there is past work that explores vibration feedback with similar parameters using conventional vibro-tactile stimulation, there are a number of factors which could have influenced the transfer from signal frequency to experienced vibration. The main two factors include:

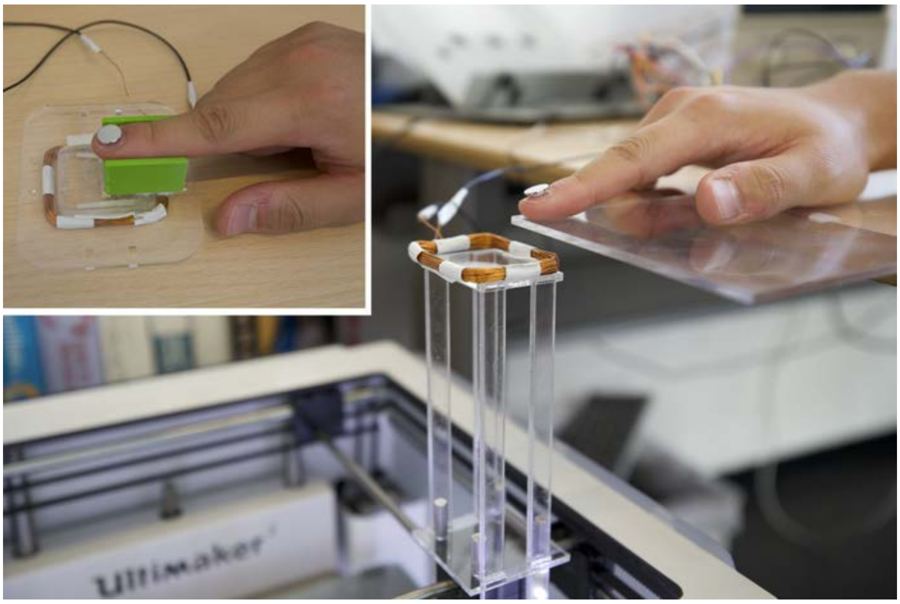

Figure 6: In both experiments, the participants finger was fixed in a single position. Top left shows the setup for the first experiment. Towards the bottom, the apparatus can be seen for moving the coil for the second experiment.

the location of the magnet on the fingernail; the transfer of electrical to electromagnetic to mechanical energy.

The haptics evaluation consists of three parts. First, to better understand how to design haptic signals, we a) present a user study of possible parameters, given constraints in tracking frequency. Using the parameter combination that lead to the strongest feedback in study ’a’, we then b) present a user study that investigates how the finger’s position relative to the coil effects the strength of the feedback. Finally, to better understand variations in strength discovered in study ’b’, we c) present measures of the electromagnetic field produced when providing feedback.

2.6.1 *a) Feedback Parameters*

We conducted an experiment to better understand what parameter combination provides the strongest haptic sensation. We recruited 10 participants. We affixed a magnet to the fingernail of the index finger of the dominant hand of each participant with double sided adhesive pads. For this experiment, we fixed the position of the finger (and thus magnet) with respect to the coil as illustrated in Figure 6, inlay. Participants were asked to rate how strongly they perceived haptic signals which varied in duration and frequency.

INDEPENDENT VARIABLE: SIGNAL DURATION The maximum duration of each signal is limited by the desired tracking frequency.

As the electromagnetic field required for providing haptic feedback interferes with the magnetic tracking, tracking and feedback must be alternated. We found that we need to delay readings by 3 ms after generating the electromagnetic field to ensure error free sampling of the magnet position. Measuring the magnets position takes 1 ms. Tracking at frequencies common with current touch solutions (60 Hz or above) provides us with a total time window of at most 16.7 ms, which supports a maximum feedback duration of 12.7 ms (per the results of the section 2.5.2).

INDEPENDENT VARIABLE: PULSE LENGTH The perception of a haptic signal is strongly linked to its frequency content [12]. To explore how the frequency effects perceived strength, we vary the length of electromagnetic pulses that each signal consists of. Signal Duration and Pulse Length interact: the Pulse Length cannot be longer than the Signal Duration, short Pulse Lengths enable signals consisting of multiple pulses, which might create stronger haptic signals. The Pulse Length is the reciprocal of its frequency.

COMBINATIONS To better understand the parameter space, we testes 11 combinations of Signal Duration and Pulse Length (12 ms duration: 2, 3, 4, 6 & 12 ms pulse length, 6ms duration: 2, 3 & 6 ms pulse length, 4 ms duration: 2 & 4 ms pulse length, 2 ms duration: 2 ms pulse length)

DEPENDANT VARIABLE & PROCEDURE For each combination, participants rated how strongly they perceived the haptic feedback. Following psychophysics methodology suggested by Gescheider [60], we asked participants to freely assign values to the feedback strength. These ratings not only provide us with a guideline on choosing a Signal Duration / Pulse Length combination, but also allow us to create visual response scales for each independent variable, as done by Strohmeier and Hornbaek [170]. Participants rated each combination three times. Combination order was randomised.

RESULTS User ratings were normalized so that the maximum value per user was 1 and ratings were then averaged per user and condition using the geometric mean. We subtracted each user's average rating from all values, so as to move all users data in the same frame of reference, where zero represents the average rating, positive values are above average and negative values below average, adapted from Strohmeier and Hornbaek [170]. The resulting scales are unit free, but the confidence intervals shown in Figures 8 and 7 provide guidance concerning size of effects.

We found that overall the combination of 12 ms signal duration with 4 ms pulses yielded the highest perceived strength. This result agrees with the literature, 4 ms pulses produce a 250 Hz signal, which is in the

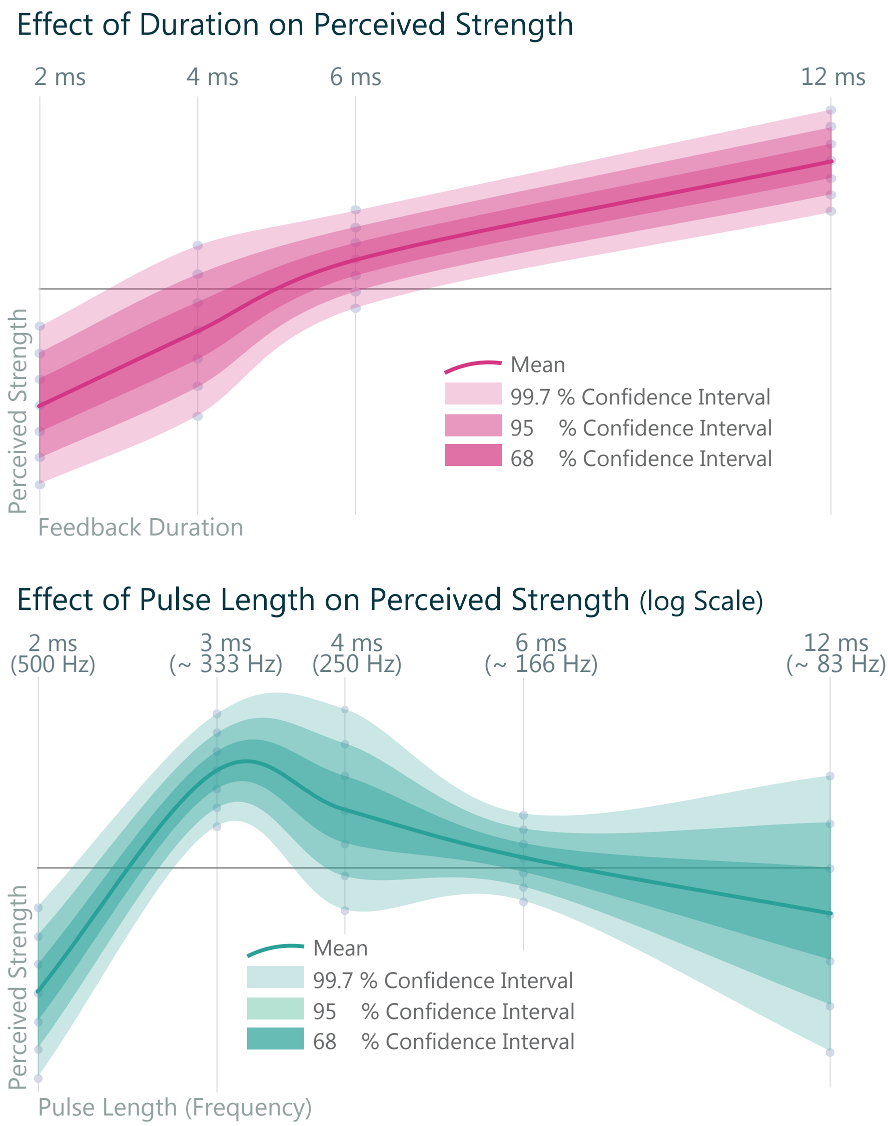

Figure 7: Perceived strength of feedback based on duration (top) and pulse length (bottom).

centre of the frequency range Pacinian cells respond to [84], stimulating the same cells responsible for texture perception [12]. The combination of 12 ms signal duration and 4 ms pulse duration supports tracking speeds up to 83 Hz. If the tracking speed is increased to 166 Hz or above (6 ms duration or above) the strength of feedback that can be provided with a discrete signal becomes significantly weaker. Mean results and corresponding 95% confidence intervals for all combinations are shown in Figure 8.

To better understand the effects of Duration and Pulse Length on their own, we also plot them individually. Figure 7a shows what appears to be a relatively linear relation between the signal's length and it's perceived strength for signal durations $<$12 ms. Figure 7b shows that

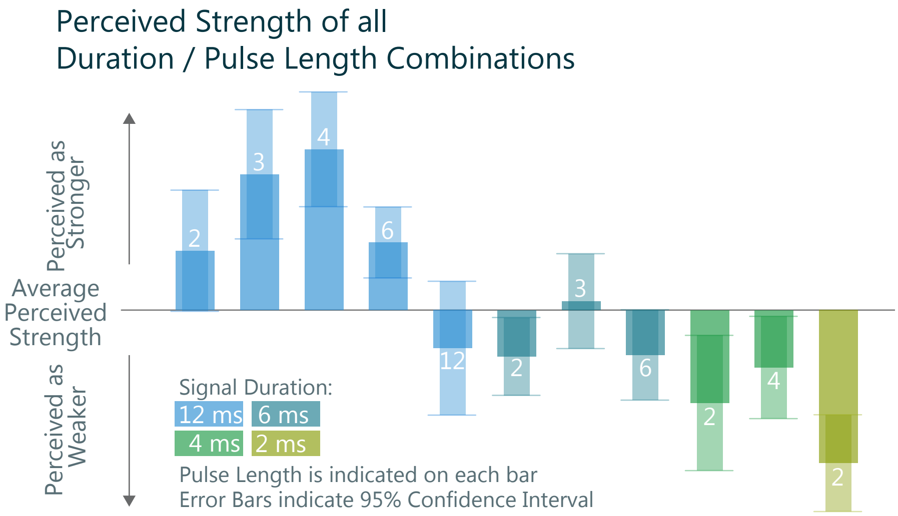

Figure 8: Our implementation only supported certain combinations of pulse length and duration. The perceived strength of each combination is shown here.

the intensity levels we found match those recorded using direct nerve readings of pacinian cells [151]. It should be noted that not all pulse lengths are combined with all signal durations, the graphs only show the combinations described previously. This is reflected in the wider confidence intervals for short durations and for the 4 ms and 12 ms pulses.

2.6.2 *b) Actuation Volume*

The second experiment was designed to explore the uniformity of the haptic feedback based on where the tracked point is positioned relative to the coil, as well as establish an ideal range within which to provide haptic feedback. Based on the result of the previous study, we conducted our exploration of the actuation volume using a feedback duration of 12 ms with a pulse duration of 4 ms.

APPARATUS The apparatus for this experiment can be seen in figure 6, which uses a modified 3D printer, that moves the coil in 3 dimensions. The coil is held above the printer by acrylic to avoid magnetic field interference that may be present nearer to the bed of the printer. As with the previous experiment, the participant is asked to keep their finger in a fixed location throughout. The experiment moved the coil into 60 different positions around the device, in a 60x60x60 mm cube. The area inside the coil was not sampled.

Figure 9 illustrates the volume that was sampled, from a Z height that starts with the coil level with the fingernail. At each location the participant was asked to rate how strongly they perceived the strength

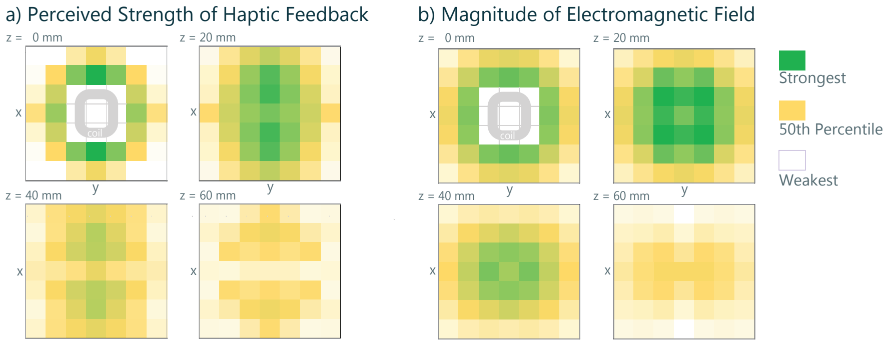

Figure 9: Left: Strength of Haptic Feedback as evaluated by participants. Right: The magnetic field as measured with a magnetometer.

of the feedback. During the experiment, the participants were blindfolded and wore headphones to avoid any visual or audial biases. Each of the 60 locations were repeated 3 times.

Since the field strength is symmetric about 3 axes, we decided to only sample an eighth of the field. The data is extrapolated in Figure 9 to mirror the x and y axes.

VARIABLES The independent variable of this experiment is the position of the finger in 3D space. The origin of this position is in the centre of the coil. We use the euclidean distance from the fingertip to the centre of the coil as a measurement to compare the strength against. Our dependant variable is the perceived strength of the feedback, as in previous experiment.

RESULTS We processed the strength ratings the same way as we did for the previous experiment. We found that, up to 4cm, the ratings were relatively similar. Beyond that, the mean results aligned with $strength = 1/distance^3$ (R = .993), which is the behavior one might expect if the feedback strength were directly proportional to magnetic field strength. Figure 10, shows the perceived strength plotted against the euclidean distance of the location to the centre of the coil.

We had much dispute over this. Personally I believe that even though the correlation is strong, it might be completely coincidental.

To understand the range that can reasonably be used for haptic feedback, we analyzed the differences in strength ratings along the distance axis. We assume that, moving from centre of the coil outward, users can clearly feel the haptic signal, as long as there is a significant difference to a weaker signal further out. We found that 56.6 mm was the last value to be significantly different from weaker signals further out (Bonferroni corrected, $p < .01$ for 84.9, 87.2, 93.8 & 103.9). While individual users might experience the signal beyond that, 56.6 mm is the furthest out that users could clearly distinguish the feedback from even weaker signals. The full confusion matrix can be found in the supplementary material.

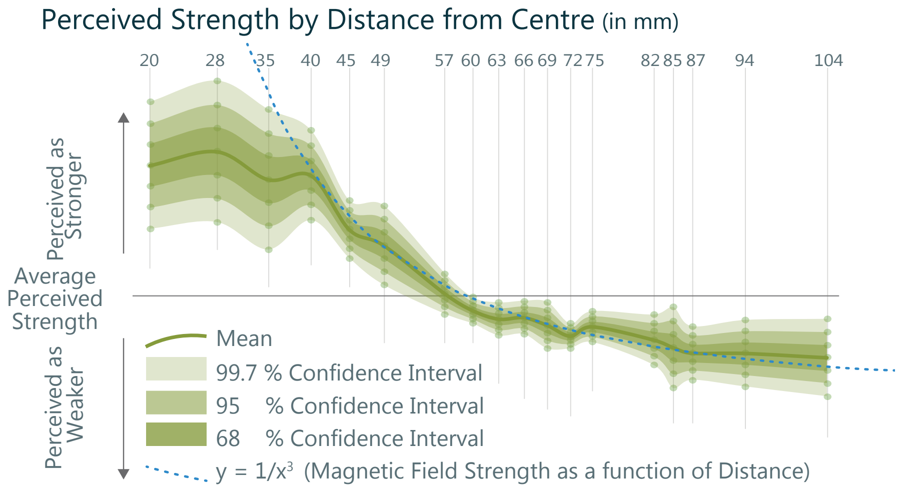

Figure 10: Perceived Strength of Feedback sorted by distance to centre.

To better understand how the positioning of the finger relative to the magnetic field effected the ratings, we created a heat-map figure 9, centre. The heat map demonstrates that, while the strength generally declines with distance, the pattern is more complex. For example, at z = 60 mm the area directly on top of the centre of the coil provides relatively weak feedback, compared above or below the centre on the x-axis. With knowledge of such patterns, haptic widgets can be created that are even further away from the device than 56.6 mm, if one places them in areas where the feedback can be clearly perceived. Alternatively, these variances in perceived strength could be corrected for, by taking the position of the finger into account when providing feedback.

2.6.3 *c) Measurements of Magnetic Field*

If the perceived strength correlated directly with the magnetic field, we could leverage the existing know-how of magnetic fields for strength calibration or widget-placement strategies, as discussed above. To better understand the patterns found in the previous user study, we therefore repeated the procedure. This time, however, instead of asking participants to rate the strength of the vibration, we placed a magnetometer where the participants finger would be and measured the strength of the magnetic field produced by the coil.

RESULTS The magnitude of the field strength correlates well with the haptic feedback ($R = .805$, Figure 9, right) but does not completely explain all variations. At z = 60 mm, the magnitude hardly correlated with the perceived strength at all ($R = -0.01$). Instead the perceived strength correlated well with the strength measured with the y-axis of the magnetometer ($R = .84$). Looking at the measured dimensions of the magnetic field individually demonstrates that they contribute to the haptic experience differently: the y-dimension of the magnetic field measures appeared most strongly to influence perception (R = .721), followed by z (R = .62). The x dimension did not correlate well with perceived strength of the feedback (R = .08). Plots for the individual dimensions can be found in the supplementary material.

2.6.4 *Summary of Findings*

To sum up, we found that users could feel signals best at a frequency of 250 to 333 Hz. We also found that within the short durations tested, there was a linear relationship between perceived strength and signal duration. The longest possible signal duration (12 ms) pairs nicely with 4 ms pulses (250 Hz), which was the combination users rated as strongest. Using this the 12 ms, 250 Hz combination, we explored the volume around the coil and found that users could reliably detect the feedback to at least 56.6 mm. We also found that there were non-linear effects, based on the relative position of the finger to the centre of the coil. We found that this pattern correlated with the magnetic field strength to some extent (R = .84), but the exact relation requires further exploration.

2.7 APPLICATIONS

See also Chapter 4 for how Magnetips *might be used with an implanted magnet.*

See also Chapter 4 for interaction methods using Magnetips and an implanted magnet.

Magnetips can track and provide feedback in three dimensions around the device, and works regardless of visual occlusions. This enables a range of novel interaction scenarios. To demonstrate some of these opportunities, we built example applications that uses Magnetips capabilities, as shown in figure 11 and in the video figure. All the applications that we have built use haptic feedback in mid-air, which is the main benefit that Magnetips brings to around-device interaction.

2.7.1 *3D Mid-air Interaction*

In past literature, there has been work that demonstrates tracking of permanent magnets even in 3 dimensions. However, by adding haptics, this enables virtual elements in mid air to be feelable.

To demonstrate 3-dimensional interaction, we created a clock application (figure 11 a). This application lets the user choose the time by rotating the finger in a radial movement around the device, similar to Abracadabra [70]. In this version, we let the user choose between minutes and hours by varying the depth of the radial movement. In addition, each step through a unit of time (second or minute) triggers a haptic signal. In this example, we demonstrate above device tracking, side of device tracking, and mid-air feedback.

Figure 11: Magnetips enables these interactions: a) 3D mid-air, b) Behind arm, c) Back of device, d) Through material.

2.7.2 *Behind Arm*

Magnetips enables interactions through occlusions made by the body itself. We prototyped an application for panning a map in 2D, by dragging the map underneath the arm (figure 11 b). Such an interaction modality may be useful when visual attention is still required for the task, as this does not occlude the display of the wearable device [11]. This interaction is also an example of skin input, which Magnetips can complement with localised haptic feedback.

2.7.3 *Back of Device*

Back of device interaction is a common research topic in mobile device interaction [8], for smartphones in particular. We built a photo editing application to demonstrate that Magnetips also enables this interaction paradigm. In the application, while the user is preparing a photo for sharing on their smartphone, the user could press and hold the 'filters' button and then use their index finger behind-the-device to scroll through the available list, receiving haptic feedback at every item boundary. This allows the user to also get visual feedback of the effect of the filter, without losing any on-screen space to displaying the list. Similarly, the technique could be used to create mid-air triggers with feedback behind the device for gaming.

2.7.4 *Eyes-free Interaction & Interacting Through Materials*

Magnetips can support eyes-free interaction in different ways. Current devices typically use vibration as eyes-free feedback channel. This causes the whole device to vibrate and can cause a buzzing sound when in contact with surfaces. Using Magnetips, the user can receive feedback off-device, directly onto their finger - supporting subtle interaction. A similar interaction can also work through clothes, for example, through a sleeve for a smartwatch or through trouser pockets, bags and purses for a smartphone.

2.8 LIMITATIONS AND FUTURE WORK

Magnetips enables reliable haptic sensations up to 56 mm from the center of the coil. At that range, the tracking error is 6.38 mm. This enables Magnetips to create a larger interaction space around mobile devices. Magnetips presents a range of new opportunities, which we explored above with our applications. There are, however, also limitations to Magnetips, which remain future work.

2.8.1 *Coil Design and Power Usage*

We designed Magnetips to use a large coil that can easily in-case existing devices. In part, this was inspired by the Qi charging coils already found in many devices.

We used a single coil for providing feedback around devices. In future work, different sizes of coils, tessellations of coils, and their impact on the haptic experience could be explored. For example, a phone-sized Magnetips device could consist of (a) one large coil, spanning the majority of its size, (b) four smaller, equally sized coils, or (c) a tessellation of different sizes. By creating different designs of coil arrays, designers could experiment with different granularities of feedback in certain locations around the device. For example, an additional small coil at the bottom of a device, could be used to drive a stronger sensation above the location of the traditional ‘home’ button.

While Magnetips has been designed to fit within mobile devices, we have yet to optimise its power requirements. The power supply in the studies was set to use a maximum of $20\,\text{V} \times 4\,\text{A} = 80\,\text{W}$, for 12 ms per tick of feedback. This equates to less than 1 Joule of energy (0.9576 Joules). This is much more than current vibration motors, for example, which would require ~0.02 Joules for the equivalent feedback. (For reference, the iPhone 7 can store 39,600 joules). The power consumption may be reduced by using an array of smaller coils. As the position of the magnet is known, one can then generate a directed magnetic field towards the position of the magnet, similar to the electromagnet array in FingerFlux [194]. This would allow the same strength of haptic feedback at a fraction of the power.

In the studies we used a power supply rather than battery to remove any variation in energy use due to the battery discharging over the course of the study. However, a mobile version of the haptics system is feasible, with a 7.4V 600 mAh 25C battery, as shown in the video figure and in Figure 12.

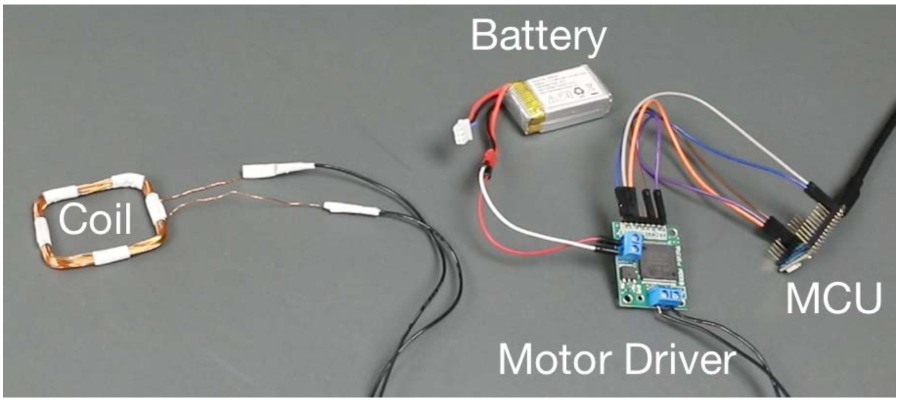

Figure 12: The haptics system can be operated with very few components, requiring only a coil, motor driver, battery and a microcontroller.

Magnetips is limited to the use of a single magnet. With passive magnets, the magnets cannot be individually enabled and disabled, making discerning individual magnets' locations an outstanding challenge. This may be solvable with an optimiser, but we are currently unaware of an existing solution. Similarly, addressing multiple magnets individually with haptic feedback is also difficult. However, a coil array may be able to provide coarse addressability at the very least.

2.8.2 *Case Design Constraints*

A ferromagnetic metal may produce an opposing magnetic field if it is shaped into a closed conducting loop, per Lenz's law. Therefore two device design parameters should be taken into account. The first of which is the material. For smartwatch or smartphone cases, ferromagnetic metals are not an ideal choice for durability or weight. Additionally, such metals interfere with magnetometer readings in the device necessary for compass data. This is why aluminium or non-metal materials are usually used in case design. The second design parameter is to avoid closed conducting loops in the case. Fulfilling either of these design constraints minimises reductions in haptic feedback intensity.

Anecdotally, when designing the example applications, the strength of the magnetic field did not change noticeably.

2.8.3 *Magnet Shape*

Users are required to wear a magnet. While this can be considered a drawback, MEMS magnetometer sensors are becoming increasingly better in sensitivity. This means that with future improvements to technology, we would be able to use smaller magnets to achieve the same level of tracking performance, making it easier and more viable to embed them into nail art [87], or even into the finger.

Experimenting with the device, we used a flat, cylindrical neodymium magnet. It is possible to acquire magnets in different shapes and sizes, including arc-shaped segments of magnets[2]. Such a shape would fit the shape of the fingernail far better. Aside from wearing the magnet, it is possible to embed magnets into interactive tools, such as a stylus, to track and provide feedback in mid-air.

2.9 CONCLUSION

With Magnetips, we have extended the interaction space to include output and well as input. Through this, Magnetips improves on previous work by not only preventing occlusion of small displays, but also requiring users to glance at the display less, as users can receive haptic guidance when interacting with the device. Magnetips also allows users to treat their skin as a supporting surface, providing on-skin input and output. Finally, Magnetips does not require line-of sight between the device and the tracked point, allowing for interaction in locations previously not possible with haptic feedback.

Through a series of technological evaluations and user studies, we showed that the current implementation can track the users finger with an average error of < 6.4 mm during a mobile task within a volume of 60x60x60 mm. Our results show that users can reliably feel the haptic feedback up to at least 56 mm. We show that we can track the users finger and provide concurrent haptic feedback that is clearly perceiveable at over 60 Hz. We suggest that haptic signals are optimised for strength at 12 ms duration and 250 Hz. If higher tracking and feedback frequencies are required, we suggest using 6 ms signal duration at 330 Hz for up to 166 readings per second.

We presented four basic usage scenarios - aimed to demonstrate how magnetips might benefit familiar interactions. We believe, however, that the opportunities offered by the unique combination of colocated tracking and haptic feedback, in combination with the ability to interact through occluding materials, might extend far beyond these scenarios, and look forward to seeing how magnetips might be adapted in other form factors and contexts.

2 http://www.neodymium-magnet.org/Curved-Magnet-p291.html

2.10 ACKNOWLEDGEMENTS

This work was supported by the European Research Council, grant no. 648785.

3 COMPARISON TO PSYCHOPHYSICS LITERATURE

3.1 MECHANORECEPTORS OVERVIEW

There are four main types of mechanoreceptors in our hands [84], of those the *Pacinian Corpuscles* are typically associated with texture perception [12] (see also Table 2). Using direct nerve readings on cats, Sato [149] demonstrated that the sensitivity of Pacinian cells increases linearly between 40 Hz and 200 Hz. The sensitivity peaks between ∼250 Hz and ∼300 Hz and then decreases again. Verillo [187, 188] demonstrates that the sensitivity of humans to vibration follows the same pattern (as indicated in Figure 13 in gray).

	Fast Adapting	Slow Adapting
Small Receptive Field	*Meissner Corpuscles* (Sensitive to ∼5 - ∼50 Hz, edges and contours)	*Merkell Endings* (Sensitive to ∼<5 Hz, static force)
Large Receptive Field	*Pacinian Corpuscles* (Sensitive to ∼40 Hz-∼400 Hz, insensitive to static force)	*Ruffini Corpuscles* (Sensitive to static force, tension)

Table 2: Comparison of receptors relevant to this work. The size of the receptive field typically relates to how deep the receptor is located in the skin – the closer the cell is to the surface, the smaller the corresponding receptive field. Fast adapting cells react to changes in stimulus, while slow adapting cells react to the presence of a stimulus.

Meissner Corpuscles are similar in structure to Pacinian Corpuscles, but sit closer to the surface of the skin and are typically associated with perception of fine surface features, edges and contures. They are responsive to lower frequency vibration from ∼5 Hz to ∼50 Hz (sometimes referred to as flutter-vibration [178]).

Meissner and Pacinian Corpuscles are both fast adapting - they respond rapidly to changes in stimulus, but do not continuously fire if a stimulus is sustained.

Slow adapting mechanoreceptors - *Merkell Endings* and *Ruffini Corpuscles* react to sustained signals. Generally speaking, they are responsive to pressure and tension, respectively. Table 2 shows an overview over some of the mechanoreceptors found in humans.

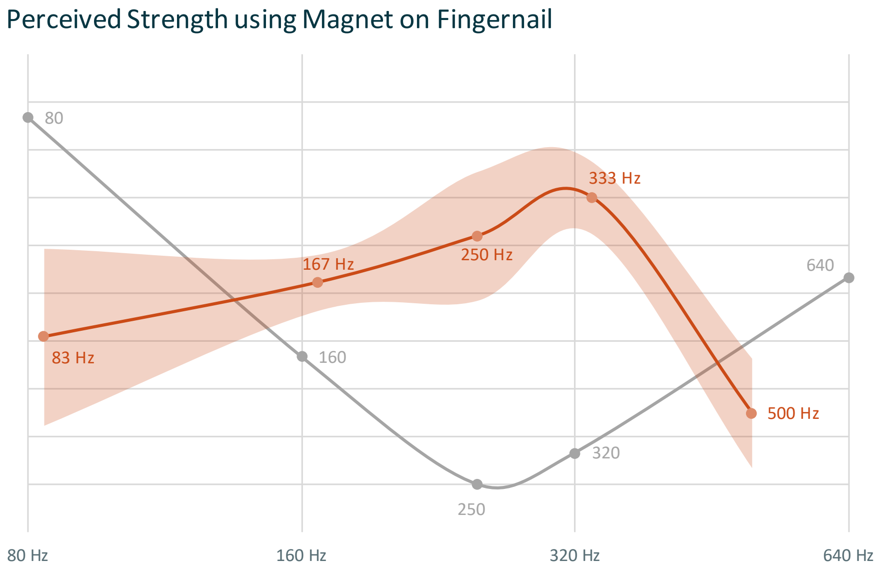

Figure 13: Mean magnitude estimation (red line) and 95% confidence interval (shaded area) of perceived magnitude as a function of frequency using the Magnetips system [115]. Perception thresholds established by Verillo [187] indicated for reference in gray. Intuitively we would expect a negative correlation between perception threshold and magnitude of sensation.

Textbooks and other introductions to sensory systems often falsely distinguish between mechanoreceptors and proprioceptors. The distinction between these two systems is not at all clear. For example, Pacinian Corpuscles are found throughout the body, including bone periosteum and joint capsules [56] and they play an important roll in our proprioception [67].

3.2 CONTEXTUALIZING MAGNETIPS RESULTS

The magnitude estimation results obtained in Chapter 2 appear to suggest that the vibration created by magnetips is mediated through Pacinian Corpuscles. Figure 13 shows both the results from our study as well as the results of a detection threshold experiment conducted by Verillo [187]. Verillo found that the minimum amplitude required for participants to perceive a haptic actuator declined between 80 Hz and 250 Hz, suggesting that, within that range, sensitivity increases as the frequency increases. Our magnitude estimation study showed that participants perceived the strength of the vibration to increase between 83 and 250 Hz. Between 333 Hz and 500 Hz the perceived strength of the stimulation provided by the magnetips system decreased. This corresponds with the increased detection threshhold between 320 Hz and 640 Hz observed by Verillo. There appears to be a discrepancy in the region between 250 Hz and $\sim$ 320/333 Hz. This could be due to measurement error or idiosyncrasies of the Magnetips system. It should be noted that a result which correlates perfectly with the data by Verillo [187] would be compatible with the result that we found (all results within the shaded area of Figure 13 are not significantly different from our result at $p < 0.05$).

As Verillo [188] later demonstrated that his measures correspond closely to those obtained by Sato [149], it is reasonable to assume that the vibrations created using Magnetips are mediated by Pacinian Corpuscles. Further study regarding the discrepancy in the peak sensitivity is however required.

4 USING MAGNETIPS WITH AN IMPLANTED MAGNET

4.1 INTRODUCTION

As described in the preface, I have a magnet in the palm of my hand. I had it implanted by a body modification artist eight years ago and it has become a somewhat mundane part of my everyday life. I had it implanted because alternating electromagnetic fields cause the magnet to vibrate. This gives some objects an extra physical dimension, for example I can feel the activity of security system at the entrance to my local library.

The system presented in *Magnetips: Combining Fingertip Tracking and Haptic Feedback for Around-Device Interaction* [115] (Chapter 2), was in part inspired by my desire to build an interactive systems around the magnetic implant. The actuation mechanism used for magnetips was copied from previous prototypes I had built to vibrate my implanted magnet in a controlled fashion. My interest in the tracking aspect of Magnetips came from the idea that it could be used to do local on-body position tracking of devices relative to the implant's location.

This section will describe some potential interaction methods using my implanted magnet and the Magnetips system, as well as present a brief exploration of the experience of the vibrotactile feedback provided to me using the Magnetips system.

4.2 INTERACTIONS

As described in Chapter 2 [115], Magnetips enables tracking of a magnet attached to the user's fingernail, in close proximity to a device. Additionally, we can provide vibrotactile feedback to that magnet, by creating electromagnetic pulses. Instead of using the system to interact with a magnet attached to a fingernail, the system can also be used to interact with an implanted magnet. Depending on how the Magnetips system is attached to the body, different types of interactions become possible.

4.2.1 *Magnetips as a loose bracelet-style device*

If the wearable device is loosely fit around the body, the magnet can be used as an anchor-point for detecting local on-body position. This can enable a user to change the device's position or orientation as an input-method. For example sliding it up or down the wrist, or rotating it around the wrist (see Figure 14). The advantage of this method over an IMU which would be the most obvious alternative) is that all motion registered would be in the local frame of the body, supporting interaction while moving, or in arbitrary poses. Additionally, translation (for example moving a bracelet towards or away from the wrist) would be hard to precisely capture with an IMU. In Figure 14 a user switches between applications by moving a smartwatch on their wrist. The Magnetips system detects that the distance to the implanted magnet increases and switches applications accordingly.

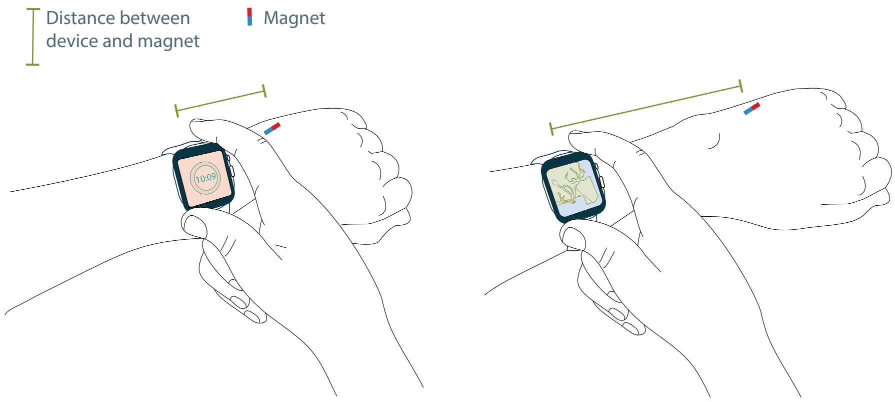

Figure 14: We can infer the position of the band by its distance from the implanted magnet. Applications can be mapped to positions on the arm. In this case, the user is switching between a time-telling application and a map.

4.2.2 *Magnetips as a tight arm-band-style device*

If the wearable device is in a known position, any changes in the magnet's position can be attributed to movements of the magnet. This in turn can be used to support gestural input. In Figure 15 the user switches between applications by changing the angle of their wrist. Here the magnetips system detects that the direction of the magnet has changed and switches applications accordingly.

As magnetips is also equipped with an IMU, a device might additionally leverage its global motion as additional cue for more complex gestural input, combining arm movements, wrist movements and device displacement. The on-board IMU can allow magnetips to intelligently switch between using the magnet position for gesture detection or local on-body position detection.

I use a smartwatch as a familiar example. Personally, I would enjoy seeing the Magnetips system embedded in garments or jewelry instead.

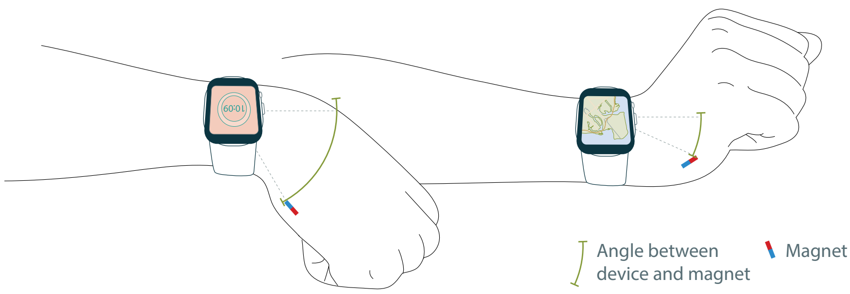

Figure 15: If the smartwatch is in a fixed position, any change in direction of the magnet can be attributed to hand movement. This enables robust gestural interaction. Here the user has applications mapped to hand-poses. We see the user switching from time-telling to map application.

Maybe even used as guidance and interaction system for epidermal robots roaming the body [42].

4.3 HAPTIC EXPERIENCES

We repeated the magnitude estimation task of Chapter 2 with myself as sole participant, using my implanted magnet. We found similar effects of duration as before, but the effects of frequency were markedly different. The experienced strength of the vibration declined above 333 Hz, corresponding both to the previous experiment [115] (Chapter 2) and with what we would expect based on detection thresholds [187]. However, below 250 Hz the perceived intensity increased, which is unlike what we previously observed or what we would have expected (See Figure 16).

As this was unexpected, we did some informal testing of a wider parameter range and different pulse durations. We found that, as the pulse duration increased, varying the frequency created qualitatively distinct experiences. While describing these beyond that they are *distinct* is difficult, some of these experiences felt sharp, while others felt less like vibration and more as if something were pulling or pushing. To see if the unexpected effect of frequency was consistent, and to gain some insight on the qualitative experience we repeated the experiment. This time we tested a wider range of frequency levels (16 Hz, 32 Hz, 56 Hz, 88 Hz, 128 Hz, 176 Hz, 232 Hz, 296 Hz, 368 Hz, 448 Hz) but only a single level of pulse duration (1200 ms). We chose the frequency levels so that the difference between the stimuli were a geometric series, and so they covered the range we were most interested in. Stimulus order was randomized and each frequency was presented four times. The entire procedure was repeated three times, once for perceived strength, once for perceived sharpness and once for perceived deformation (experience of pulling or pushing).

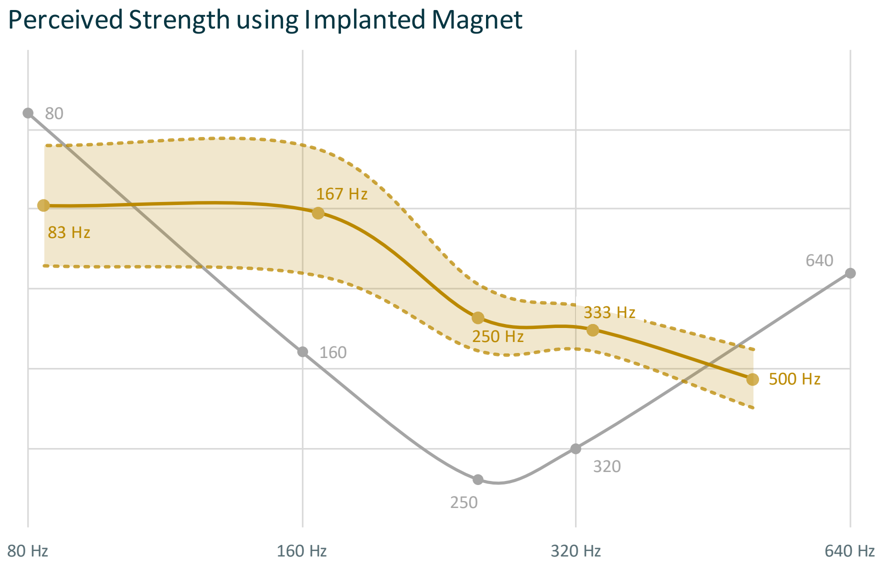

Figure 16: In-vivo vibrations: Estimates of experienced strength of vibration, using the same settings as for the Magnetips evaluation (See Figure 13 for reference). Solid line is my mean response, dotted lines are two standard deviations (roughly equivalent to a 95% confidence interval). Human sensitivity to vibration added in light gray for reference.

The results can be seen in Figure 17. We found that, again, the perceived strength increased below 250 Hz, peaking at 88 Hz and then decreasing gradually (Figure 17, top). This is surprising compared to the results of the Magnetips study [115] and to what one would expect based on detection threshold alone [187], but consistant with our initial experiment using the implant (Figure 16).

The experience of sharpness (Figure 17, middle) had a slight positive correlation with frequency. There is a dip in perceived strength at 128 Hz, which, on its own, might appear to be noise. However, it correlates with a peak in sharpness and a dip in deformation. The 176 Hz level displays the opposite behavior. This odd behavior is, at the very least, a repeatable effect with the setup we used.

The experience of deformation (Figure 17, bottom) was strongest at low frequencies (~100 Hz and lower) and decreased at higher frequencies.

The qualitatively different experience below ~100 Hz (high deformation, low sharpness) compared to above 100 Hz (low deformation, high sharpness) suggests that the way the vibration is mediated might be different. This observation also makes sense in light of the area below ~ 100 Hz being closer to the receptive range of the Meissner Corpuscles (see Chapter 3).

However, as the stimulation device is an experimental device, and we do not completely understand the relation between electrical pulse, magnetic field, and actual motion of the magnet, I am hesitant to draw any conclusions stronger than that the experience with the implanted magnet differs from the experience with the magnet attached to the fingernail. The strange behavior at 128 Hz and 176 Hz might also be an artefact of the magnetips implementation, however we have no specific suspicion of why this might be the case.

4.4 CONCLUSION

While the experience of stimulation provided to the implanted magnet is different to that of the magnet attached to the fingertip, the haptic feedback still can clearly be felt. In fact, the clear qualitative differences experienced as the frequency changes might allow Magnetips to provide a large number of distinct signals to the user. The tracking of the magnet is unaffected, as the body is not ferromagnetic. The interaction methods I present suggest that an implanted magnet can provide a useful anchor for local on-body positioning and that this can open up the design space to a wide number of new interactions for wearable and on-body devices.

Please do not see this as an endorsement for implanting magnets into your body. I do not have the knowledge to speak of the safety or wisdom of doing this. If you decide to implant a magnet, make sure it is coated with a bio-compatible coating. Off-the-shelf magnets are most likely toxic.

ACKNOWLEDGEMENTS

Thanks to Jess McIntosh for running these studies and thinking through the interaction scenarios with me.

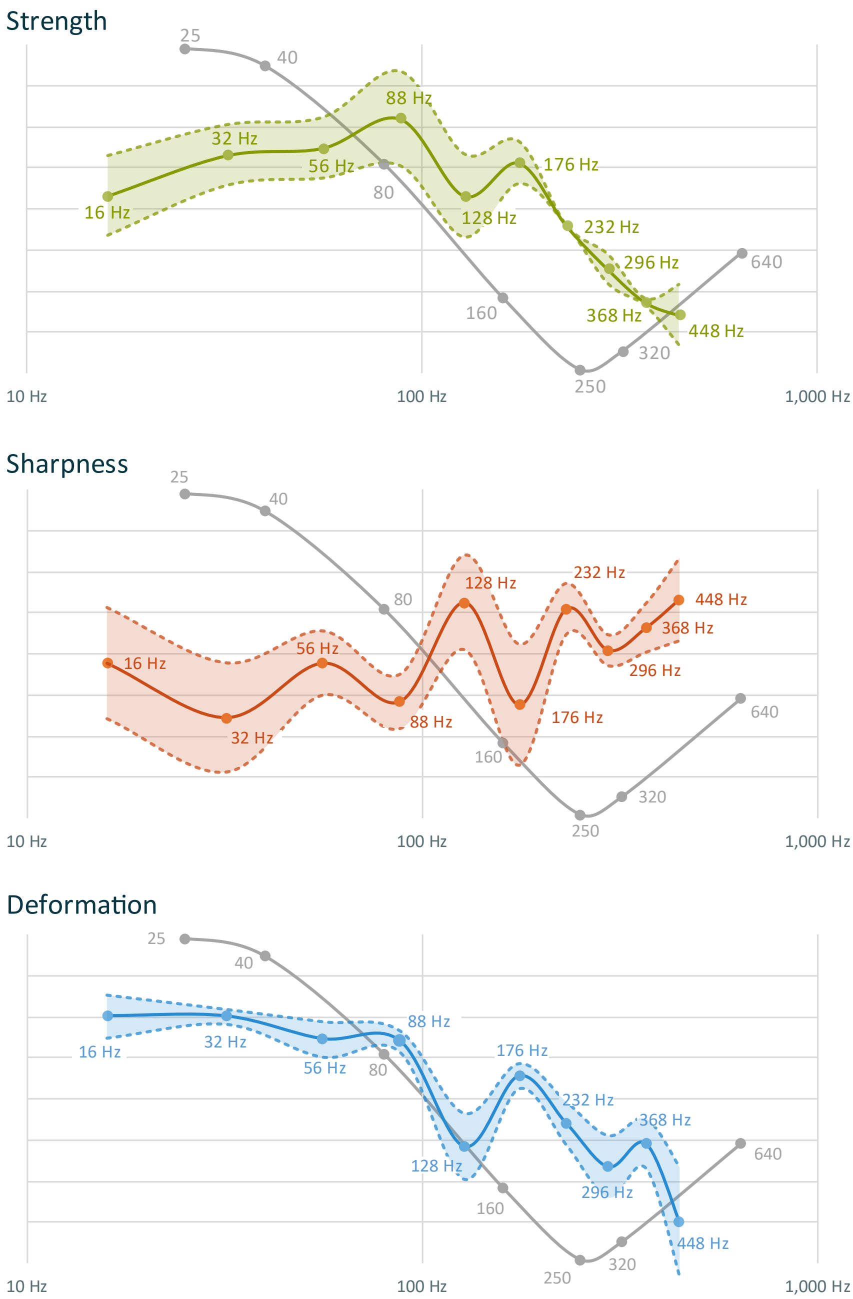

Figure 17: In-vivo vibrations: Estimates of experienced strength and sharpness of the vibrotactile stimulus and estimate of perceived deformations such as pulling or pushing. Solid line is my mean response, dotted lines are two standard deviations (roughly equivalent to a 95% confidence interval). Human sensitivity to vibration added in light gray for reference.

5 REFLEX: A FLEXIBLE SMARTPHONE WITH ACTIVE HAPTIC FEEDBACK

This work was conducted at the Human Media Lab, Queen's University, Kingston, Ontaro. It was a collaboration with Jesse Burstyn, JP Carraascal and Vincent Levesque, supervised by Roel Vertegaal.

Citation

Paul Strohmeier, Jesse Burstyn, Juan Pablo Carrascal, Vincent Levesque, and Roel Vertegaal. "ReFlex: A Flexible Smartphone with Active Haptic Feedback for Bend Input." In: *Proceedings of the Tenth International Conference on Tangible, Embedded, and Embodied Interaction.* (TEI '16). Eindhoven, Netherlands.

Abstract

ReFlex is a flexible smartphone with bend input and active haptic feedback. ReFlex's features allow the introduction of sensations such as friction or resistance. We report results from an experiment using ReFlex in a targeting task, as well as initial users' reactions to the prototype. We explore both absolute and relative tactile haptic feedback, paired with two types of bend input mappings: position-control and ratecontrol. We observed that position-controlled cursors paired well with relative bend feedback, while rate-controlled cursors paired well with absolute bend feedback to indicate targets. We also explored an eyes-free condition. Results suggest that while eyes-free, haptic feedback conditions were more error-prone than visual-only conditions, the size of the error was relatively small, and users were able to complete the task in all cases. We present two application scenarios that take advantage of the unique input and output modalities of ReFlex and discuss its potential for within document navigation.

5.1 INTRODUCTION

When presented with a tool or a device, we assess its features and attributes to understand what we can do with it, i.e, its perceived affordances [125]. Visual inspection alone cannot convey all of this information and it is often necessary to feel objects with our hands to gain a fuller understanding of its material and structural properties. We perceive these qualities both through tactile stimulation on our skin and kinesthetic receptors in our hands. When an object is deformed,

Figure 18: Reflex – a flexible, haptic smartphone.

a rich set of sensations come into play to inform us about its internal structure. For example, when reading a paper document, the physical structure of pages can help guide users to particular locations in the document. A book can have physical tabs to indicate chapters and its pages might have dog ears to denote bookmarks or creases at frequently read passages. The distribution of pages between the hands provides some haptic representation of the current reading location. And pages sliding between a user's fingers provide feedback on the speed with which she is navigating. Many, if not all, of these haptic affordances are lost when navigating documents on rigid Tablet PCs. When designing flexible organic user interfaces [76], the structural qualities of the device are inherited from the material of the substrate used in its construction. While haptic technologies have been used to mimic textures of different materials [10], these technologies often focus on surface features. An alternative approach is to modify the perceived material properties of the device [90, 206].

5.1.1 *Reflex: A Haptic Flexible Smartphone*

Contributing to the latter approach, we created ReFlex, a flexible smartphone featuring a high-resolution flexible display and a haptic actuator (Figure 18). ReFlex modifies the experience of dry friction when bending the device, as perceived through tactile and kinesthetic receptors of the fingers. This feedback creates possibilities for new interactive experiences, such as simulating the elastic and material properties that occur while navigating a paper book. We report on results from a study that explores bending using position control and rate control in combination with different types of haptic feedback in a target acquisition task. We found that position control resulted in faster times and higher accuracy than rate control. Participants preferred relative haptic feedback with position control and absolute haptic feedback with rate control. Following these results, we performed a study evaluating the effectiveness of the haptic rendering techniques for acquiring a target in the absence visual feedback. We found that it was possible to accomplish the task with a relatively low error rate. Based on user observations of the two studies, we discuss the possibilities of our haptic rendering techniques to enhance navigation for long digital documents, providing some of the same qualities as paper books. We present two application scenarios that may improve the browsing experience of lists and documents on flexible smartphones.

5.2 RELATED WORK

5.2.1 *Elastic Input Devices*

Elastic devices and corresponding control mappings were also studied by Casiez and Vogel [31]

Zhai [209] distinguishes between two types of devices: isometric and isotonic. While their suitability for different types of cursor control has been extensively studied, their defining feature is their stiffness, i.e. how much they oppose physical displacement. Isotonic devices, e.g., a mouse, have a constant low resistance and are freely moved. Isometric devices, such as, e.g., the IBM TrackPoint [148], fully resist displacement and operate through forces applied. Between them, however, is a third category: that of elastic devices. These have a stiffness, k, has a resistance that is proportional to its displacement. Elastic devices signal their displacement through passive force feedback, and, like springs, are naturally self-centering. Many flexible display devices fit into this category since when deformed, they flex and return to their original state upon release. Changing an elastic device's stiffness moves it along the isometric-isotonic spectrum. Higher amounts of resistance can afford better rate control, while lower resistance affords more proprioceptive feedback during displacement and can be more suitable for position control [209].

5.2.2 *Bendable Devices*

Flexible display interactions is a relatively recent, but increasingly popular, field of study. Early explorations, such as Gummi [155], predate the actual use of flexible displays. Some of the first explorations using real flexible displays include Lahey et al.'s PaperPhone [100] and Nokia's Kinetic [91]. Along with these prototypes, there have been several studies of how users perceive the physical properties of deformable devices. Nakagawa et al. [122] presented MimicTile, a bendable device with dynamical stiffness. They demonstrated that participants could accurately identify different levels of stiffness and Kildal et al. [92] reported that users preferred flexible devices that are less stiff than others.

5.2.3 *Haptics and Perception*

I did not use the words perception and experience as carefully when this paper was written as I do now. The text here is kept as published.

There are many studies on how people use haptic cues to infer an object's properties; most are well beyond the scope of this paper. For example, tactile and kinesthetic cues can be used to create an illusion of texture. Klatzky et al. [94] outline models for a force feedback mouse that simulated varying levels of surface roughness. With TeslaTouch, Bau et al. [10] use electro-vibration to create dynamic friction on a touch surface. Lederman and Jones [102] present a literature survey on how manipulating sensory cues can create both kinesthetic and tactile illusions. Changing visual [23] or auditory cues [19] can create false perceptions of an object's stiffness, while changing an objects configuration can create varying perception of its weight [74]. Conversely, changing the haptic sensory cues can also change the perception of material stiffness [206].

5.2.4 *Haptics and Performance*

Researchers have found that haptic and tactile feedback can benefit pointing tasks when used to provide direct information about the target. In a Fitts' Law targeting task, Akamatsu et al. [4] found that tactile feedback provided on the target resulted in equivalent movement times, but shorter final positioning times, than visual or audio feedback. Forlines et al. [52] reported that haptic signals on targets are beneficial for both crossing and pointing Fitts' law tasks. On the other hand, Kildal et al. [92] demonstrated that passive haptic feedback, in the form of device stiffness, had little to no effect on task performance, but greatly influenced user comfort and feedback on bend interaction.

5.3 DESIGN RATIONALE

We were interested in exploring the interaction between bend input with passive force feedback and actuated vibro-tactile feedback. Specifically, we wanted to understand what it means to combine variations of haptic feedback with different styles of bend interaction.

5.3.1 *Passive Haptic Feedback*

When using a touch screen, kinesthetic feedback provided by the configuration of the arm is largely independent of touch location. Compared to such traditional touch interaction, the ReFlex's passive force feedback provides a strong coupling of proprioceptive feedback with bend input: ReFlex provides a linear correspondence between the applied force and position or speed of a cursor.

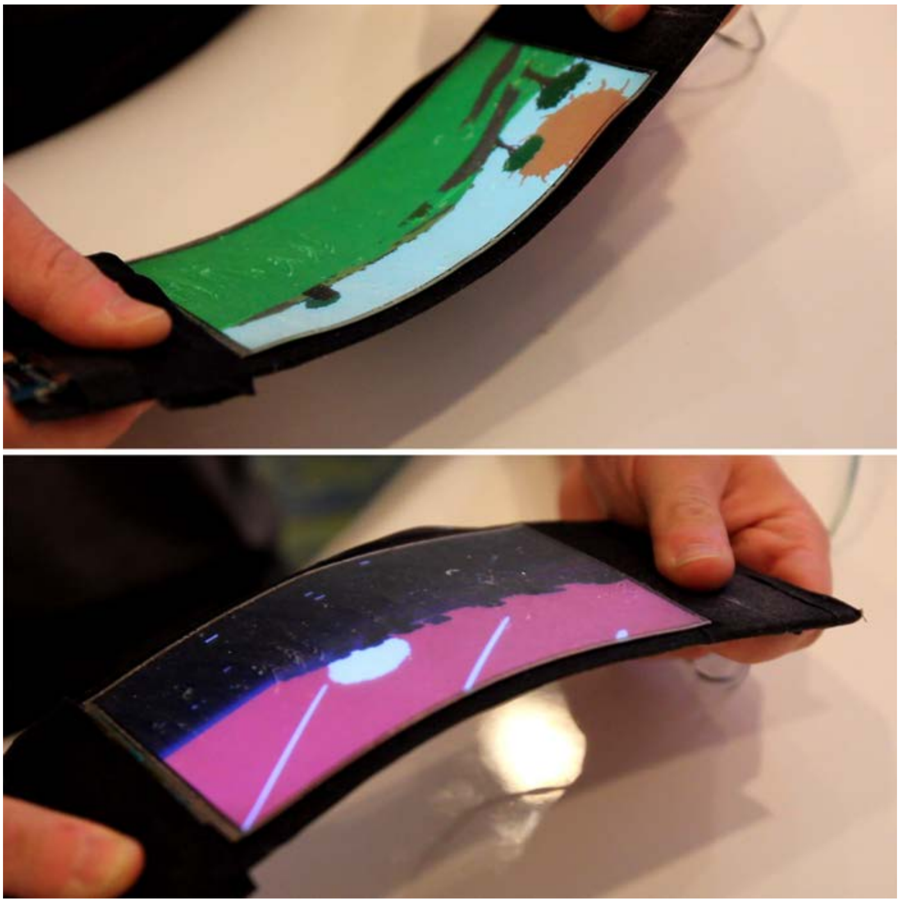

Figure 19: ReFlex bending in both directions.

5.3.2 *Adding Active Haptics to a Bendable Smartphone*

We facilitate multisensory feedback and explore the interplay between visual, tactile and kinesthetic experiences. To accomplish this, we augmented the flexible display with a haptic actuator that provides active feedback in addition to the passive elastic forces generated by the device when bent. This enables us to actively modulate the experience of passive haptic feedback experienced by bending the device, generating variations in the perceived elasticity and internal structure.

5.4 IMPLEMENTATION

ReFlex is a flexible smartphone prototype with a bend sensor and haptic actuator (Figure 18). Our prototype can be used as a stand-alone device and runs Android 4.2 (Figure 19)

5.4.1 *Display*

Reflex uses a FOLED display manufactured by LG Display. The 6.0" (135 mm × 77 mm) FOLED display has a resolution of 1280 × 720 pixels and a refresh rate of 60 Hz. The display is mounted on a flexible substrate that extends 5 cm to left and right. This allows for a comfortable grip without occluding the display, and a location for rigid electronics, such as the processor and display driver board.

5.4.2 *Rigidity*

The structural substrate of the device is designed so that it is most flexible at its center, tapering off towards the semi-rigid ends. This was accomplished by adding trapezoidal reenforcements on both ends of the device (Figure 20).

5.4.3 *Input*

ReFlex measures the direction and extent of a bend with an Omega Engineering strain gauge [131] placed at its center. A Teensy 3.1 microcontroller samples this strain gauge at 12 bits and ∼2000 times per second. The high temporal and spatial resolution of this strain gauge allows us to synchronize the haptic actuation with the passive haptics that are naturally present when bending the device. ReFlex also has a multi-touch sensor which we use for setting up the experiments, and a button on the back, which participants use during the experiment.

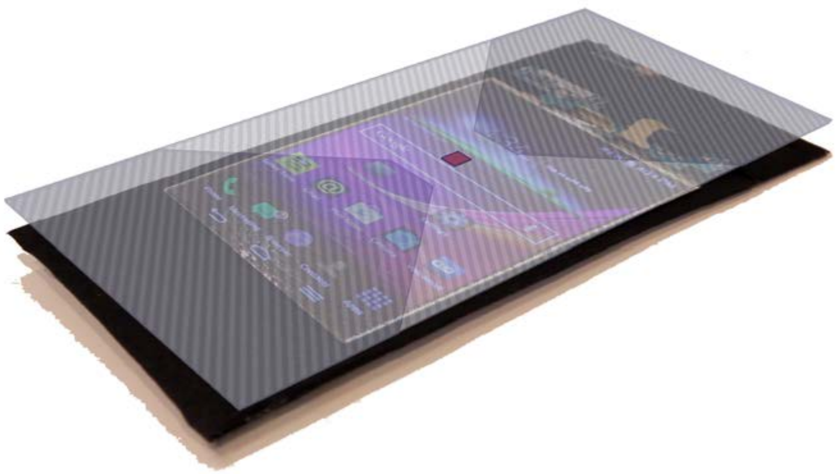

Figure 20: Prototype with internal structure: Dark gray areas are semi-rigid, light gray areas are flexible. Position and size of strain sensor is indicated by the red square.

5.4.4 *Haptic Output*

Active haptic feedback is generated using a Tactile Labs Haptuator [177] mounted on the back of ReFlex, parallel to the length of the display. ReFlex uses a vibrotactile transducer, rather than a vibrational motor, for precise temporal control of haptic signals. The Haptuator generates precise discrete pulses at an acceleration of up to 73 m/s^2 and at rates up to 1000 Hz [177]. Discrete pulses are inaudible, and a continuous series of pulses are audible only above ~500 Hz. The Haptuator is meant to be driven by a standard audio amplifier. For our study, it was driven by the sound card of the computer running the experimental software.

5.4.5 *Software*

A C++ program polls the Teensy microcontroller for sensor values and converts them into cursor movement 200 times per second. It then passes these values to a Max patch that generates audio signals for the Haptuator. ReFlex runs a simple Android client application that receives cursor and target information from the computer over WiFi to draw them on the display. ReFlex is tethered to optimize the synchronizations between haptic and visual feedback, however it can also be used as a stand-alone, wireless device.

5.5 BEND INPUT MAPPINGS

See also Chapter 6 for an alternative explanation of mapping types.

ReFlex uses two different types of cursor control, following the definitions put forward by Zhai [209]:

POSITION CONTROL (PC) For ReFlex, position control means that sensor values are directly mapped to pixels on the display's x-axis, i.e., the cursor position (p_c) is linearly proportional to the amount of bend (x) applied to the device:

$$p_c(x) = x$$

When the device is flat, the cursor is in the center of the display. The cursor is at the left extremity of the display when ReFlex is fully convex, and on the right extremity when ReFlex is fully concave. This mapping creates a linear correspondence between passive kinesthetic feedback and the visual position of the cursor on the display.

RATE CONTROL (RC) With rate control, the user controls position of the cursor by manipulating its speed and direction via bend gestures. The direction of movement is same as in Position Control. The speed at which the cursor moves (r_c) is mapped to the extent of the bend with the following sinusoidal easing function:

$$r_c(x) = -\cos\left(x \cdot \frac{\pi}{2}\right) + 1, \quad x : [0,1]$$

5.6 HAPTIC FEEDBACK ALGORITHMS

Active tactile haptic feedback was created using an audio signal consisting of a train of discrete pulses. We applied a high-pass filter (at 80 Hz) to the signal to remove any low frequency elements, helping each pulse to become more distinct. A low-pass filter (at 200 Hz) helped to attenuate the audibility of the signal. We also modulated the amplitude of the signal so the haptic pulses would be felt stronger at the extremes of bend input (∼160 db) and weaker when the device is close to rest (∼110 db). This increased the perceived strain proportionally to the how much the device is bent. We used two types of bend input to haptic feedback mappings, as explained next. Depending on the mapping, we could create the haptic illusion of altered material properties, similar to 3D-press [90].

5.6.1 *Haptic Feedback Mapping Types*

See also Chapter 6 for an alternative explanation of feedback types.

ABSOLUTE BEND FEEDBACK (A) In this mapping, the rate of the pulse train (r_p) varies linearly with the extent of bend (x). The more extreme the bend, the higher the pulse train rate:

$$r_p(x) = x$$

The duty cycle of the pulse train is 50 % for all rates. When combining this mapping with a cursor rate control (A-RC), the pulse rate varies with the speed of the cursor. This created the feeling that fixed locations on the display trigger haptic pulses when the cursor passes over them. On the other hand, when combining this mapping with cursor position control (A-PC), the synchronization between cursor location and haptic pulses is lost. The device simply pulses faster the further it is bent.

RELATIVE BEND FEEDBACK (R) for this mapping, the pulse train is not necessarily periodic. Instead, its rate varies linearly with the speed of the bend movement, i.e., the bend velocity:

$$r_p(x) = \frac{\Delta x}{\Delta t}$$

Each single pulse of the train is a 1ms length square pulse. When combining this mapping with cursor position control (R-PC), absolute locations on the display seem to trigger haptic pulses as the cursor passes over them. Conversely, when using rate control (R-RC), the pulse rate seems to be synchronized with the acceleration of the cursor.

5.7 EXPERIMENT 1 - EFFECTS OF FEEDBACK AND MAPPING ON TARGETING PERFORMANCE

Participants performed a subset of a one-dimensional Fitts' Law, targeting task. Two vertical ribbons appeared on the display, with varying center-to-center distances. Target width was held constant at 80 pixels. Users were asked to alternately click within the left and right ribbon 25 times. Each block of trials began after the participant placed the cursor within the left target and pressed the button. Participants were instructed to perform the task as quickly and as accurately as possible.

80 pixels corresponds approximatly to 77mm.

5.7.1 *Experiment Design*

We used a 3 × 2× 3 factorial within-subject design with repeated measures. Our factors were haptic feedback (3 levels, discussed below), cursor control (position control and rate control), and target distance (150, 500, 960 pixels). Participants performed one block of 25 trials for each of the 18 combinations of factors. Condition order was counterbalanced between participants. Participants practiced with each combination of haptic feedback and cursor control until they achieved less than 10% improvement between trials. Our measures were targeting time and error rates.

To avoid learning effects, participants practiced before the experiment until they felt comfortable using the device using both cursor control mappings. Additionally, they actually performed 29 trials of which we discarded the first 4.

5.7.2 *Haptic Feedback*

Participants were provided with 3 levels of haptic feedback: no feedback, absolute bend feedback, and relative bend feedback. Depending on the cursor control, the active feedback provided information on cursor position (A-PC), cursor speed (R-PC, A-RC), or cursor acceleration (R-RC).

5.7.3 *Participants*

12 participants performed this experiment (7 male, 5 female) with ages ranging from 20 to 38 years. Most participants (10/12) were right handed.

5.7.4 *Experiment 1 Results*

We analyzed targeting times using a repeated measures ANOVA on haptic feedback (3) × cursor control (2) × target distance (3). The analysis showed that cursor control was a significant factor ($F_{1,11} = 251.02$, $p < .001$), with position control resulting in faster targeting times than rate control. We also found that target distance was a significant factor ($F_{2,22} = 339.01$, $p < .001$) and there was a significant interaction between cursor control and target distance ($F_{2,22} = 108.40$, $p < .001$). We found no significant effects of either type of haptic feedback. We analyzed the errors using a repeated measures ANOVA on haptic feedback (3) × cursor control (2) × target distance (3). The analysis showed that cursor control was a significant factor ($F_{1,18} = 9.86$, $p < .05$), with position control resulting in fewer errors than rate control. We also found that there was a significant interaction effect between cursor control and target distance ($F_{2,36} = 0.74$, $p < .05$).

This is actually super interesting: Even though the haptic feedback provided additional cues to the user's actions, this did not result in a measurable difference.

5.8 EXPERIMENT 1 DISCUSSION

5.8.1 *Rate Control vs. Position Control*

Users were able to complete the task faster using position control than using rate control. When first comparing rate control to position control most users also commented that they did not like rate control. After using the rate controlled input for a longer period of time, however, participants commented they also found rate control easy to use. Some users stated that rate control was preferable for targets at large distances, while position control was better for targets at short distances. This was, however, not supported by the targeting times; in fact, the interaction effect we observed indicates the opposite. The error rates appeared constant for all target distances in rate control, while for position control larger target distances caused more errors. This may in part explain participants' experiences: while position control is both faster and more precise, the precision benefit over rate control is greatest for short movements.

We wondered if physically changing the stiffness, rather than virtually changing the perception of the device, might lead to a different result. Burstyn at al. [26] followed up this work with a more in depth analysis, this time using devices with variable stiffness. They reproduced the effects of input method, and they did not find any significant effect of device stiffness on performance.

HAPTIC RENDERING We did not find any measurable effects of the different types of haptic feedback on task completion times. Interestingly, this contradicts the feedback comments we obtained from participants: many of them found the haptic feedback helpful. This result does not contradict previous research that observed effects of haptic feedback on pointing tasks only when direct information about the target is signaled (e.g. crossing the center or edge) [4, 52], as in our experiment we only provided indirect target information such as speed and position. We examined participants' comments to understand how the haptic feedback might have influenced their opinion on the prototype. P2 stated that "if I specifically pay attention to the feedback, it is helpful, [but I think] my brain responds faster to my eyes than to my finger", suggesting dominance of the visual system in the task [4]. P6 commented that "the haptic feedback is helpful; it allows me to focus less on the visuals." We found this comment interesting given the measured results; despite the fact there were no differences in performance, the participants felt it was the case. It is possible that the haptic feedback can generate an additional perceptual illusion, in the same vein to perception of animation speed in moving progress bars [71].

5.8.2 *Combinations of Haptic Feedback and Cursor Control*

See also table 5 in Chapter 6.

The two different types of cursor control types and methods of generating haptic feedback can be combined in four unique ways. Each combination results in a very different experience for the user:

R-PC:
RELATIVE BEND FEEDBACK,
POSITION-CONTROLLED CURSOR As the cursor moves across the screen, the user receives pulses of haptic feedback when crossing fixed locations on the screen. When the device is released, the cursor selfcenters as a result of the elasticity of the device and position control. The haptic pulses become stronger towards the edge of the device; the active haptics correlate with the user's perception of passive bend forces. These properties seem to make this combination one of the easiest to interpret. For instance, P1 stated "I can almost close my eyes and roughly have an idea of where the cursor is", an idea that we explored further in Experiment 2. The synchronization of feedback with body and display shape makes this haptic configuration unlike regular vibration. Instead, it is experienced as friction within the device, creating rich haptic images. Participants explained it in diverse ways: "It almost feels lit it's more fibrous" (P1), "It feels like I'm bending a twig of wood" (P3), "It's a little bit like when you are moving a rubber band along a smooth surface" (P4).

A-PC:
ABSOLUTE BEND FEEDBACK,
POSITION-CONTROLLED CURSOR The cursor behaves the same way as for the previous combination, however, the pulses are no longer synchronized to specific locations. Instead, the further the display is bent, the faster the pulse rate. This gave an experience somewhat reminiscent of flicking pages of a book at a rate that corresponds to the exerted force. P4 explains it is "because when you bend it and hold statically the vibrations just continue at a steady pace", while P3 states "that's not a trait I would attribute to an inanimate type of object that I am working with". While the mapping is not direct, participants did experience this combination potentially useful. P6 considered that "it does a very good job in providing the user with some sort of feedback as to the amount of pressure to exert on the screen", while P5 considered it to be "consistent throughout", and P3 suggested that this combination "could be something very useful in gaming".

A-RC:
ABSOLUTE BEND FEEDBACK,
RATE-CONTROLLED CURSOR In this cursor control, the amount and direction of bend determines the speed and direction of the cursor. The haptic pulses appear at regular intervals which become shorter the further the device is bent. This provides an effect of the pulses being triggered by the cursor moving over fixed locations on the display; the haptic pulses are experienced as the texture of the surface the cursor is passing over, or as explained by P1, "It feels like it's bumping a regular number of times as it moves across the distance (. . .) It's like a texture, like a gradient". This combination of cursor control and feedback style was experienced as intuitive. "The haptic feedback just makes me more aware of the speed of the cursor" (P3); "It helps me know when to stop moving the cursor or when to slow it down" (P4).

R-RC:
RELATIVE BEND FEEDBACK,
RATE-CONTROLLED CURSOR With this cursor control, the feedback has haptic pulses triggered with changes in cursor speed. It appeared difficult to interpret and makes the interaction feel disjoint: "I'm not entire sure of what the haptic feedback is indicating" (P3); "It doesn't particularly feel like the feedback is helping me" (P4); "[the task with this haptic feedback is] not necessarily easier, it supplements but doesn't make it easier" (P5). P1 and P6 suggested, "I don't like this".

With combinations R-PC and A-RC, the cursor movement and active haptic output appear closely synced, leading to a predictable behavior of the haptic qualities of the device. For combination R-PC, it is experienced like the internal structure of the device is changing, while A-RC is experienced much like modulating surface textures that the cursor moves over. While A-PC did not have these properties, its clear mapping was still considered useful. R-RC did not have the same type of coupling experienced with R-PC and A-RC, nor did it have a clear mapping, like combination A-PC. Unlike all other combinations, participants did not enjoy it. With these results in mind, we used the optimal combinations to test if the prototype would allow a user to perform a task in an eyes-free scenario.

5.9 EXPERIMENT 2 - EYES-FREE HAPTIC TARGETING

We conducted a second experiment to assess the effects of haptic feedback for indicating targets without visuals. Participants performed the same targeting task as in Experiment 1, with the same apparatus but a somewhat larger target width of 120 pixels. We encouraged the participants to prioritize accuracy over speed.

5.9.1 *Haptic Feedback*

The participants used haptic feedback to find the target. In one condition, haptic feedback only occurred when the cursor was over the target. In another, the haptic feedback was removed only when the cursor was over the target. A distinction from other investigations that utilize haptic feedback to signify targets [4, 52] is that our feedback is based both on the presence of a target, as well as the behavior of the cursor. That is, in both haptic conditions, the user only feels haptic feedback when the cursor is moving. We chose the feedback types that participants felt closely matched cursor control: we used relative bend feedback with a position-controlled cursor (R-PC) and absolute bend feedback combined with the rate-controlled cursor (A-RC).

5.9.2 *Experiment Design*

We used a 3 × 2 × 2 factorial within-subjects design with repeated measures. Our factors were feedback method (visual feedback, haptics present only when on target, and haptics always present except when on target), cursor control (position control and rate control), and target distance (500 and 960 pixels). Our dependent measures were number of errors and error size.

5.9.3 *Participants and Training*

6 participants performed this experiment (4 male, 2 female), with ages between 20-26. Most (5/6) were right handed. They were given 90 seconds to explore each combination of haptic rendering.

5.9.4 *Experiment 2 Results*

We analyzed errors using a repeated measures ANOVA on feedback method (3) × cursor control (2) × target distance (2). The analysis showed that feedback method had a significant effect ($F_{2,10} = 7.20$, $p < .05$) on the number of errors. Post-hoc tests, with Bonferroni corrected comparisons, revealed that visual feedback had significantly fewer errors than the condition where haptic texture was removed from the targets. There were no significant effects of either cursor control or target distance on error rates. Table 3 outlines the mean number of errors for each combination of feedback method and cursor control. Table 4 presents the mean error distances.

		Cursor	
		Position Control	Rate Control
Feedback Type	Visual Only	1.00 *(0.95)*	1.00 *(0.95)*
	Haptics On Target	4.17 *(6.32)*	11.00 *(8.38)*
	Haptics Off Target	3.17 *(4.71)*	5.00 *(2.95)*

Table 3: Mean number of errors in Experiment 2.

		Cursor	
		Position Control	Rate Control
Feedback Type	Visual Only	11.09 *(0.95)*	10.90 *(10.51)*
	Haptics On Target	23.15 *(17.6)*	60.03 *(39.69)*
	Haptics Off Target	22.94 *(23.99)*	31.91 *(28.74)*

Table 4: Mean error distance (pixels) in Experiment 2.

5.10 EXPERIMENT 2 DISCUSSION

Unsurprisingly, the error rate and average error distance were higher in the eyes-free conditions. However, participants were able to complete this experiment without visuals in both cursor control conditions. The largest average error size was around 60 pixels (~5.8 mm), the smallest average error size was less than 25 pixels (~2.5 mm) from the target. Participants were split between their preference of haptics on targets vs. haptics off targets. Some preferred the haptics on target condition, because it felt as though they were notified once they reached the target. Others preferred the haptics off targets, as they felt it was more continuous. Transitioning from the in-between space to the targets during the haptics off target condition was described as an interesting sensation. Participants used colorful descriptions in trying to capture this experience. P2 stated that it was as though "the space between the targets is land and the haptic feedback is water. It's like falling into water". Other descriptions used were "It's like sand and ice". Often, haptic feedback areas were referred to as "coarse", and the blank ones as "smooth."

5.11 AFFORDANCES OF REFLEX

The rich metaphors expressed by the participants indicate that ReFlex is capable of eliciting haptic sensations beyond what we would expect of traditional vibro-tactile feedback. The interplay between the passive force feedback and active tactile feedback of the device provides perceived physical affordances that can be controlled to match the requirements of the task at hand. These affordances can emulate mechanisms that we are familiar with from the physical world. For example, in the same way as we assess the length of a book by bending it and flipping through its pages, a shorter digital document would generate less haptic pulses when the device is bent than a longer one. Frequency of pulses would simulate page-flipping speed. These methods can also provide implicit information of one's usage history: just as a physical book tends to open to a section that a reader has studied intently, we could gently guide a user to the most visited sections in a digital document by varying the perceived separation between pages – i.e., the consecutive haptic pulses. We believe that the haptic feedback methods we demonstrated are suitable for providing a user with haptic renderings of content, inspired by the physical affordances and wear and tear of physical media.

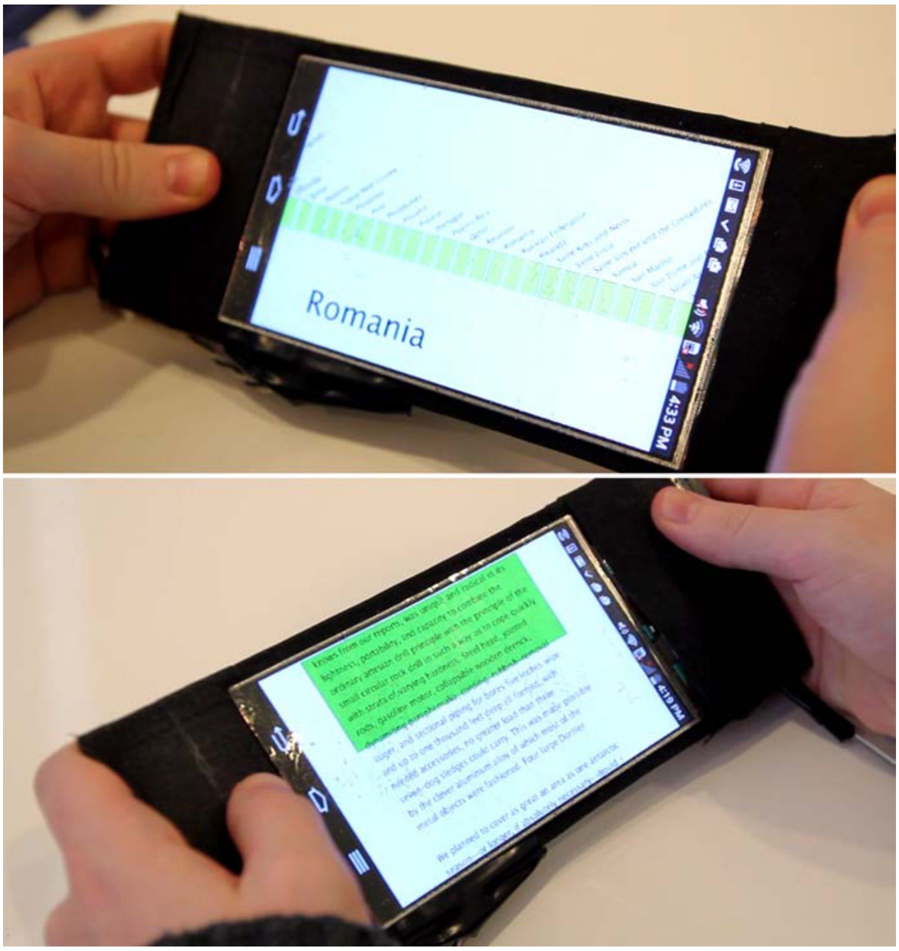

Figure 21: Top: Using A-RC for off-screen browsing of large lists, and combination R-PC for selecting on-screen items in the list. Bottom: Using R-PC for scrolling a text, and A-RC for annotating items.

5.11.1 *Haptic Qualities for Reading*

Improving the affordances of digital documents to better facilitate within-document navigation is of increasing importance, as people use more digital devices such as the iPad or Kindle to consuming magazines and books. While using these devices for accessing digital documents has numerous advantages, there are downsides to not using paper in reading tasks. These are discussed in O'Hara and Sellen's widely cited work comparing paper documents with their digital counterparts [126], as well as in Marshall and Bly's report on navigation in paper documents [113]. A critical message from these works is that the haptic affordances of paper provide users with serendipitous within-document navigation methods that are lost in reading digital documents. Previous research has proposed several ways to emulate the physical affordances of paper. Some examples include the dual slate reader presented by Chen et al. [34] and Girouard et al.'s DisplayStacks [65], based on multiple thinfilm displays, among others. Our findings suggest that flexible devices that combine active and passive haptic feedback may provide an interesting approach to improving within-document navigation for digital documents. The affordances of ReFlex can be used to support the types of serendipitous navigation that we are accustomed to from paper documents, in a form suitable for digital content.

5.12 SAMPLE APPLICATIONS

Based on these thoughts, we created two applications that take advantage of active and passive haptic feedback to support within document navigation. Like our experiments, these applications use the button on the back of the device.

5.12.1 *Large List Navigation*

Rate control and position control can act synergistically for bend based input. Precise on-screen targeting actions can use a position-controlled cursor, while off-screen actions that require fast motion or continuous input, such as scrolling, are better suited for rate control [209]. A scenario that takes advantage of this technique is navigation through large lists. Figure 21 (top) shows a user navigating a large list with bending gestures, using rate control for off-screen list browsing and position control when selecting on-screen items from the list. Users can select rate control by pressing and holding the back-of-the-device button during bends. When the button is released, ReFlex uses position control. Items are selected from the list with a click of the button in both cases. The haptic feedback switches between absolute and relative such that pulses always occur at the transition between items on the list. For item selection (R-PC) these are experienced as physical obstacles the cursor moves over, and for off-screen scrolling (A-RC) users experience obstacles as items enter the display.

5.12.2 *Text Navigation and Annotation*

Our second application is inspired by the wear and tear of physical documents that occurs while reading. Figure 21 (bottom) shows our e-reader application, which features a highlight function. Users scroll through off-screen content using rate control. Bending ReFlex into a concave shape moves the text up; the opposite moves the text down. When the user clicks the button, the application switches from reading mode to highlight mode. When in this mode, users can use position control to highlight lines of text. Highlighted areas are identifiable in two ways: visually via a brighter foreground color, as well as haptically through a texture. When scrolling, a highlighted area entering the viewport is experienced as having more friction than its surrounding text. As a user scrolls, this change in texture allows them to feel that they are passing a highlighted section – even when not attending visually or when scrolling quickly. The additional friction invites the user to pause at a previously highlighted section, like a dog ear invites the reader to open the book to a previously highlighted page.

5.13 LIMITATIONS

The preliminary evaluation presented in this paper was intended as a starting point for a more thorough study of the perceived material properties of bendable devices when augmented with haptic rendering. We obtained significant results for some variables from our experiments, but our sample size was small. We therefore do not consider it a fully conclusive or exhaustive experimental study. Our users' feedback was, however, informative and valuable.

5.14 CONCLUSION AND FUTURE WORK

We presented ReFlex, a flexible smartphone with passive and active haptic feedback. The evaluation of the prototype indicates that it has potential for enhancing document browsing tasks. We discussed this possibility and we presented two application scenarios, one for browsing long lists and another for text navigation and annotation. Many open questions remain, and we hope to address them in future work. The combination of bending, active haptic feedback, and different input-to-cursor mappings allowed us to create an extremely expressive device. By modulating several parameters of the haptic pulses – such as rate, amplitude and filtering – a rich haptic design language can be developed, one that could be used to incorporate sensations of material and structure in interface design. Further investigation is needed to empirically determine the suitability of this language for enhancing applications such as browsing digital documents.

5.15 ACKNOWLEDGEMENTS

We thank Danny Grant for his support of this project. This work was funded by grants from NSERC and Immersion Corporation.

I also thank Antonio Gomes for his support in figuring out how to actually get the flexible displays to work.

6

REFLEX REVISITED: OPPORTUNITIES AND LIMITS OF MULTI-MODALITY

ReFlex was my first serious attempt at using vibrotactile feedback for haptic rendering. ReFlex was also an attempt to demonstrate what one might be able to do with a bendable smartphone. Without realizing it at the time, these two contributions led to conflicting goals in the evaluation. We intended to demonstrate the utility of the device we built. We did this by designing targeting tasks that showed (a) how one might best control a cursor by bending the device, (b) how such control strategies might be combined with vibrotactile feedback and (c) how the haptic feedback helps complete a task eyes-free. None of these contributions systematically explore the experience of people who used the device. This was not because we were not interested in user experience, but rather because we assumed that changes in how the device is experienced would result in measurable changes in performance. This assumption did not hold.

We did find that some combinations of user control and vibrotactile feedback combined to create something qualitatively different, which we call material experience. Our own intuition and that of users who tested ReFlex in pilot studies strongly suggested that these material experiences did influence targeting behavior. To demonstrate the effects of material experiences, we designed a third experiment. Recognizing that the effects were subtle, we designed the haptic feedback in a way we assumed would clearly demonstrate their utility: we changed the material experiences to indicate that the user had reached the target.

Spoiler alert: we were wrong

6.1 RECAP: FEEDBACK AND CONTROL METHODS

For sake of clarity, I will repeat the definitions of the control and feedback methods used in Chapter 5

6.1.1 *Control*

A more detailed description can be found in Section 5.5.

POSITION CONTROL The user controls the cursor position. When the device is flat, the cursor is in the middle. When the device is convex the cursor is at the left side of the display. When the device is concave it is at the right side of the display. Each bend level has a corresponding cursor position.

RATE CONTROL The user controls the cursor speed. When the device is flat, the cursor is stationary. When the device is convex the cursor moves to the left side of the display. When the device is concave the cursor moves to the right side of the display. Each bend level has a corresponding cursor speed.

6.1.2 *Feedback*

RELATIVE FEEDBACK The pulse frequency is determined by how much the device is bent relative to its previous state. A large change in bend corresponds to high frequency, a low change in bend corresponds to low frequency. If the device is kept steady at any position, no feedback is provided. In other words *feedback is coupled to user movement.*

ABSOLUTE FEEDBACK The pulse frequency is determined by the bend position of the device. The more the device is bent, the higher the frequency of the pulses. If the device is flat, then no feedback is experienced.

6.2 WHAT IS A MATERIAL EXPERIENCE

We found that depending on how control and feedback were combined, these were experienced very differently (Table 5). When *rate control* was used with *relative feedback*, the resulting experience was difficult to understand. The vibrations felt somewhat arbitrary, and participants found them confusing and distracting. When *position control* was used with *absolute feedback*, user could make sense of the combination, and thought it was potentially useful; however, it was experienced as vibrations.

For the other two combinations, something exciting happened. The input mapping and feedback merged into new experiences. *Position control* combined with *relative feedback* was experienced as changing the material of ReFlex, while *rate control* combined with *absolute feedback* was experienced as friction of the cursor.

I call such experiences that emerge from congruent combinations of vibration and user action material experiences. This phenomenon of new experiences emerging when vibration is coupled to user action is what the rest of this thesis explores.

6.3 CAN HAPTIC FEEDBACK IMPROVE TARGETING PERFORMANCE?

Experiment one of the *ReFlex* paper (See Chapter 5) compared different types of haptic feedback, including two that lead to material experiences, however, the haptic feedback was uniform – it changed the experience of bending the devices, but it did not provide any specific information regarding the targets. To investigate if additional haptic information might influence the pointing behavior, we repeated exper-

	Pulse Frequency Based on User Motion (Relative Feedback)	Bend Position (Absolute Feedback)
User controls cursor speed (Rate Control)	(a) "I'm not entire sure of what the haptic feedback is indicating."	(b) "It feels like it's bumping a regular number of times as it moves across the distance (...) It's like a texture." (Perception Shift)
User controls cursor position (Position Control)	(c) "It almost feels like it's more fibrous." (Perception Shift)	(d) "When you bend it and hold statically, the vibrations just continue at a steady pace."

Table 5: The combinations of (b) rate control and absolute feedback as well as (c) position control and relative feedback changed the experienced material properties of the device. I also refer to this phenomenon as "perception shift".
(a) Rate control paired with relative feedback was simply confusing, while (d) position control paired with absolute haptic feedback felt as one might expect haptic feedback in a handheld device to feel.

iment one (see Section 5.7), but this time the presence or absence of haptic feedback indicated to the user that they had reached a target. As we were specifically interested in those combinations which created new material experiences, we only used the two corresponding feedback-/mapping combinations.

6.3.1 *Experiment 3: Haptically Supported Targeting*

Ten participants – who had previously participated in experiment one (see Section 5.7) – performed the same subset of a one-dimensional Fitts' Law targeting task again: Two vertical 80 pixel ribbons appeared on the display, with varying center-to-center distances (see Figure 22). Users were asked to alternately click within the left and right ribbon 25 times. Each block of trials began after the participant placed the cursor within the left target and pressed the button. Participants were instructed to perform the task as quickly and as accurately as possible.

The main difference between this experiment, and experiment 2 of the previous chapter is visual feedback. Experiment 2 had no visual feedback, while this experiment uses visual feedback as seen in Figure 22.

We constrained the overall error rate (M: 5.4% SD: 2.2%) by asking participants to repeat blocks if they made more than 4 errors and asking participants to increase their speed if we saw the error rate drop below 4%.

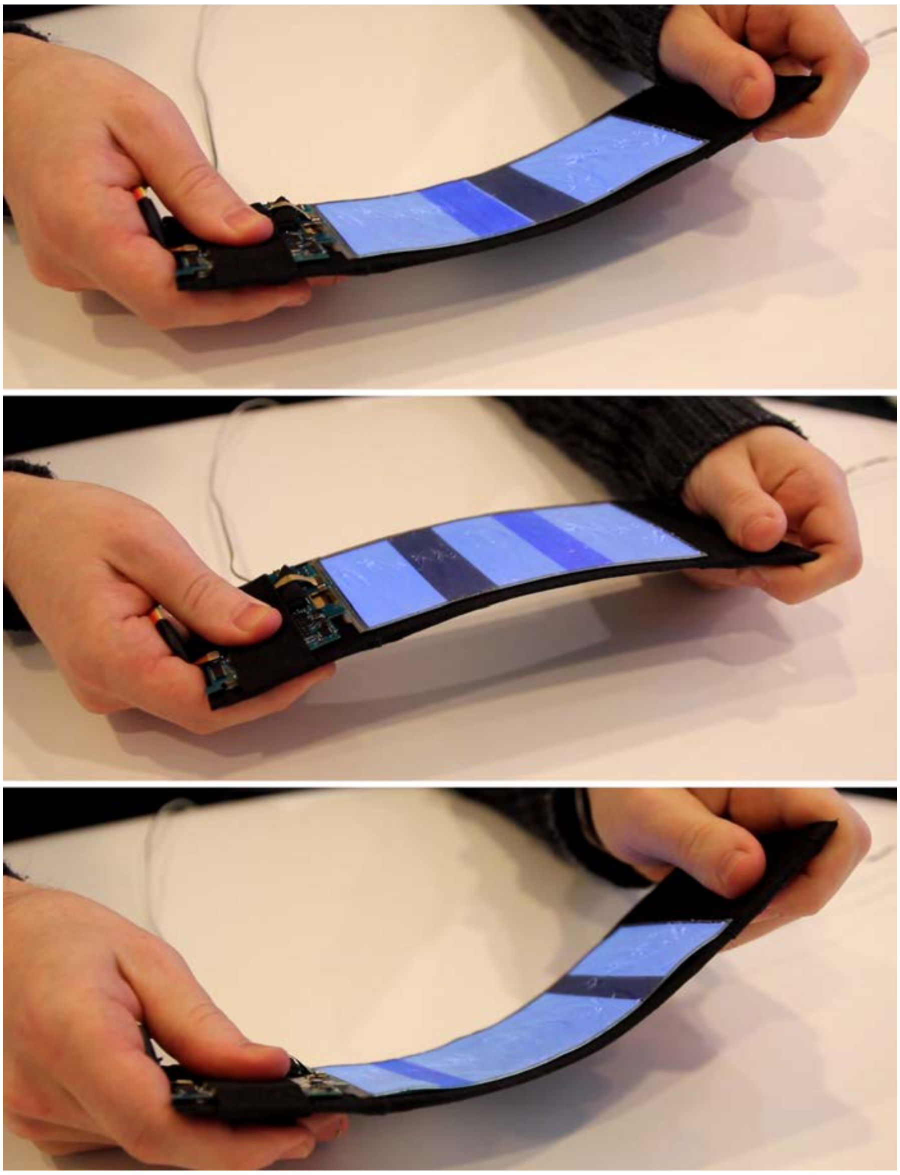

Figure 22: User performing experimental task.

6.3.2 *Conditions – Haptic Feedback*

We compared three haptic conditions (1) no haptic feedback, (2) feedback on the target, and (3) feedback everywhere except on the target. We used both combinations of feedback for which new material experiences emerged.

6.3.3 *Experiment Design*

As before, we used a 3x3x2 factorial within-subject design with repeated measures. The factors were distance (150, 500, 960 pixels), haptic feedback (3, as discussed above) and cursor control (position

Haptics	Average Targeting Times	95% Confidence Interval Lower Bound	Upper Bound
None	1150.22	920.07	1380.38
On Target	1128.26	1008.96	1247.57
Off Target	1091.78	931.68	1251.88

Table 6: Targeting times per haptic condition.

Haptics	Average Errors	95% Confidence Interval Lower Bound	Upper Bound
None	6.1%	4.4%	8.3%
On Target	4.7%	2.6%	6.9%
Off Target	5.4%	3.5%	7.3%

Table 7: Recorded Errors per Haptic Condition

control and rate control). Participants performed one block of 25 trials for each of the 18 combinations of factors. As before, condition order was counterbalanced among participants, and participants practiced with each combination of haptic feedback and cursor control until they achieved less than 10% improvement between trials. Measures targeted time and error rates.

The variability of the non-haptic conditions is consistently larger than of the conditions using haptics. This is reflected in the confidence intervals.

6.3.4 *Analysis*

We conducted an ANOVA on the targeting times of the first 21 successful trials using distance (3) × haptic feedback (3) × cursor control (2) × repetition (21) as factors. Similarly to experiment one of Chapter 5, we found that target distance ($F_{2,12} = 126.02$, $p < .001$, $\eta_p^2 = .955$) and cursor control ($F_{1,6} = 295.132$, $p < .001$, $\eta_p^2 = .98$) had significant effects. They also had a significant interaction effect ($F_{2,12} = 12.6$, $p < .001$, $\eta_p^2 = .783$).

Even though the average targeting times decreased for haptic feedback (See Table 6) the result was not significant ($F_{2,12} = .933$, $p = .421$, $\eta_p^2 = .134$). We also analyzed the repetitions to see if there were any additional learning effects due to haptic feedback, but did not find any

6.3.5 *Experiment 3: Errors*

We constrained the error rates to ensure that changes in difficulty would result in changes in targeting times, rather than changes in errors. However, as we did not find any effects of haptics on targeting times, we wondered if this might have been unsuccessful. A preliminary glance at the data (Table 7) suggests that haptic feedback reduces errors. However, after conducting an ANOVA to analyze errors, we found the effect was not significant.

We analyzed the errors per condition. We conducted an ANOVA on distance (3) × haptic feedback (3) × cursor control (2). We found no effect of distance ($F_{2,18} = 1.067$, $p = .365$, $\eta_p^2 = .106$) or haptics ($F_{2,18} = 2.422$, $p = .09$, $\eta_p^2 = .234$). We did find a slight effect of mappings ($F_{1,9} = 6.395$, $p = .032$, $\eta_p^2 = .415$). It should be pointed out that we constrained the possible error range, which would limit our ability to detect significant effects. A slightly different study design would be required to better understand the effects of haptic textures on precision.

6.3.6 *Discussion*

In experiment one of Chapter 5, we tested two conditions where the material experience of the interaction was modified by haptic feedback. In one condition users experienced additional haptic cues while bending, in the other condition users could feel friction as the cursor moved over the screen. Even though participants felt as though this helped them complete the tasks, this was not reflected in the data.

In this chapter, we conducted the same experiment again, but this time we only changed the material experience to indicate that the user is entering or exciting a target. We assumed that – if there was an effect of changing the material experience – this experiment would surely capture it. However, again, the effects we found were small and not significant.

In Chapter 5, we report on an experiment very similar to the one presented in this Chapter. The difference was that, in Chapter 5 we provided no visual information to the users. The users had to find and click the targets focussing only on the changes in material experience. Here we argue that, because users were able to complete the task, there clearly is an effect of material experience.

It appears that if a task – such as a Fitts' Law-style targeting task – is presented to the user with sufficient visual information to complete it with maximum efficiency, adding information on other sensory channels does not further improve the performance. I assume that, if tasks provide users with successively less visual information, the effects of the haptic feedback would similarly become stronger.

The trends found in our data suggest haptic feedback may have very small effects – even in this very visual task – but that these effect sizes are too small to be found in the relatively small samples sizes of our experiments and that they are not relevant for the type of task we are interested in.

Experimental Device	Demo Device
Strain Gage (Omega Engineering)	Resistive Flex Sensors (Flexpoint)
Sampling at 12 bit (∼ 4000 usable values)	Sampling at 8 bit (∼ 700 usable values)
Sampling 200 times per second	Sampling 60 times per second

Table 8: Comparison between sensing setup of experimental and demo device.

6.4 latency and multi-modality

We demonstrated ReFlex at TEI 2016 in Eindhoven. As the demonstration device was a different version of the device than our initial implementation which we used for our experiments, we were able to gather additional anecdotal observations. The two main differences between the devices were sensing fidelity and applications. The demo device had a scaled-down bend-sensing system compared to the experimental device (see Table 8) and the system was demonstrated using two applications designed and implemented by Jesse Burstyn, Ze Ye and Roel Vertegaal (see Figure 23).

The first application was an interactive comic book which users could flip through by bending the phone. The second was an Angry Birds clone, where users could catapult birds by bending the phone. The comic book app used *rate control* together with *absolute feedback* for simulating page-flipping. The Angry Birds clone used *position control* with *relative feedback*. Additionally, a varying bandpass filter was applied which had a high centre frequency when the phone was strongly bent and a low centre frequency when it was only bent slightly. This provided an effect of pulling elastic, as one does when catapulting angry birds.

Video of the demo-session at TEI can be viewed here: https://youtu.be/nh47viOTo-Y.

The demo was met with great enthusiasm. Subjectively, we felt that the reduced sampling rate and fidelity reduced the impact of the material experience. However, people trying out the applications – experts and novices alike – all provided us with extremely positive feedback. Experts in haptic feedback design told us that it was especially the merging of the visual and haptic elements which made it such a strong experience. Some visitors (typically those who had read our publication) tried the system eyes-free and commented that they felt the material experience did not work as well without the visual element. This reaction was similar to our own impression.

Figure 23: Applications using the two mappings

6.5 CONCLUSION

The experiences of demoing ReFlex live and the follow up experiment and analysis presented in this chapter provided three insights: (a) that visual information dominates tactile information when both are present, (b) that it's important to sample the user's motion with high precision, and (c) that there are only negligible or no effects of varying the material experience on pointing.

As we did not find any effects of varying the material experience on pointing, I decided to avoid evaluating material experiences with performance oriented tasks in the future. While the second study of ReFlex demonstrates that situations where there is a benefit of varying material experiences can be created, clearly there are other ways that this type of haptic feedback influences the experience of using an interactive system. Instead of viewing material experiences through the lens of performance, from now on I intend to focus on the experience itself.

The extent to which the lower fidelity of the sensing system changed the strength with which the vibrotactile signals and the user motion merged to a material experience was subjectively strong. The change in users' reactions between demo system and experimental system was noticeable. Because of this, sampling rate and fidelity will be central design considerations in future material experience explorations.

The dominance of visual information is first hinted at by the null effect of haptic feedback in a visual pointing task. It is further reinforced by the fact that the strong visuals of the demo applications appeared to mask the lower fidelity of the haptic feedback. To avoid visual information as a confounding factor, I decided to avoid using visual feedback at all in future experiments.

Part II

PERCEPTION

In this section, I investigate the experience of interacting with devices that provide haptic feedback coupled to user motion. First, in *Generating Haptic Textures*, I explore the design of the haptic signals. Using magnitude estimation, I observe how changing parameters of the haptic signal influences how people evaluate their experience. In the second paper, *Pulse Trains*, I explore how various mappings from motion to feedback are experienced. For this paper, I conducted in-depth interviews which provide rich descriptions of the process of perceiving haptic feedback from which material experiences emerge.

These papers are special to me as they were my first research papers which were not driven by a technology I wanted to develop, or by an engineering problem I hoped to solve. Instead, the research questions stood in the foreground, and any technology developed was created in service of that research question. These papers are also special to me as they were the first papers where all aspects, from ideation to execution, were under my control. These papers put me far out of my comfort zone, and I feel that I learned a great deal from them, not only about perception, but about conducting research in general.

7 GENERATING HAPTIC TEXTURES WITH A VIBROTACTILE ACTUATOR

This paper was awarded an Honourable Mention *(top 5% of all papers) at CHI 2018.*

Citation

Paul Strohmeier and Kasper Hornbæk. "Generating Haptic Tex-tures with a Vibrotactile Actuator." In: *Proceedings of the Conference on Human Factors in Computing Systems* (CHI 2018). New York, New York, USA: ACM Press

Abstract

Vibrotactile actuation is typically used to deliver buzzing sensations. But if vibrotactile actuation is tightly coupled to users' actions, it can be used to create much richer haptic experiences. It is not well understood, however, how this coupling should be done or which vibrotactile parameters create which experiences. To investigate how actuation parameters relate to haptic experiences, we built a physical slider with minimal native friction, a vibrotactile actuator and an integrated position sensor. By vibrating the slider as it is moved, we create an experience of texture between the sliding element and its track. We conducted a magnitude estimation experiment to map how granularity, amplitude and timbre relate to the experiences of roughness, adhesiveness, sharpness and bumpiness. We found that amplitude influences the strength of the perceived texture, while variations in granularity and timbre create distinct experiences. Our study underlines the importance of action in haptic perception and suggests strategies for deploying such tightly coupled feedback in everyday devices.

7.1 INTRODUCTION

Active exploration is required if one wishes to understand the texture of an object: When resting a finger on a material, we perceive the material's basic features, such as cues related to shape and temperature. To understand the texture, we also need to know what it feels like to move one's finger over it. The relative motion of the fingertip and the surface create vibrations [107] and these vibrations activate sensory receptors in our fingertips [13]. Through them the material comes alive in our hands [12].

In HCI it is becoming more common to apply this insight in the design of vibrotactile feedback, leading to systems that tightly couple vibration to human motion. Such feedback has been used to change the perceived compliance of materials [90, 171], emulate mechanically complex systems [206], and simulate contact with different surfaces through a proxy tool [40, 146].

While a haptic experience of texture is caused by our body moving relative to a surface or material, it is not clear what characteristics of a material make us experience its texture in a particular way. Consequently it is not clear how to manipulate vibrotactile feedback if one wishes to generate a specific haptic experience. Previous explorations have either used discrete mappings between pulse trains and motion or pressure [90, 171], or attempted to recreate the original sensation as closely as possible by recording the vibrotactile signature of a surface and playing it back [40, 146]. However, the first approach is difficult to generalize, as we do not understand how the experience would change if the mapping is changed, while the second approach does not contribute to an understanding of why different types of vibrotactile actuation are experienced in certain ways.

To better understand how the experience of texture can be manipulated by varying parameters of vibrotactile feedback, we conducted a magnitude estimation experiment [162]. The experiment uses haptic feedback that is tightly coupled to user input. To achieve this, we created a slider consisting of a glide-bearing that moves over an anodized aluminum rod. This bearing is augmented with a vibrotactile actuator and a position sensor with high spatial and temporal resolution. As a user moves the slider over the rod, we provide haptic pulses synchronized to the user's motion. This is experienced as a texture between the slider and its track. In the experiment, we adjusted the parameters with which we generated the vibrotactile feedback, while asking participants to rate their experience of the texture.

We found that bumpiness, roughness, adhesiveness and sharpness all had unique granularity and timbre profiles, which suggests that these parameters can be used to generate qualitatively distinct sensations. The response curve of amplitude displayed different slopes, suggesting that sensations such as bumpiness or roughness benefit from high amplitude vibration more so than adhesion or sharpness. The results also suggest that pulse frequency might play a less important role than expected and that timbre should be further investigated for a better understanding of haptic experience.

7.2 RELATED WORK

Our work is about haptic experiences, in particular, the experience of texture. A host of work exists on classifying and mapping such experiences (e.g., [203]). The vocabulary we will use is based on Okamoto et al. [130], who presented a synthesis of dimensions of haptic perception from 18 studies. Okamoto et al. identified three major perceptual dimensions: hard/soft, cold/warm and a texture dimension of rough/smooth. They also suggested that the roughness dimension has micro and macro sub dimensions, and that sticky/slippery could be another possible dimension [130]. We use roughness and bumpiness as more colloquial terms for micro and macro roughness, while we use adhesiveness to capture the sticky/slippery dimension. Sharpness was added based on user feedback during a four person pilot study. We first discuss the role of vibration in setting about such experiences. Then we survey techniques for creating vibrotactile feedback and show how coupling them to user movements help create experiences of haptic textures.

7.2.1 *The Role of Vibration in Experiencing Surfaces*

A discussion of the different types of touch recepters can also be found in Chapter 3, an overview is provided in Table 2.

To fully experience the haptic qualities of a material, touch alone is insufficient. Resting one's hand on a material may evoke an impression of temperature or reveal shape features if they are prominent enough to distort the skin, but to feel how hard a material is, one needs to actively press against it; to experience its texture, one must move ones finger relative to the object one is touching [88, 95]. When one moves a finger over a surface, the texture of the fingertip in combination with the texture of the surface produce vibrations [13, 107]. These vibrations are used to infer information about the material we are touching [12, 103]. In The World of Touch, Katz [88] differentiated between the sensation of vibration and that of pressure. He argued that either can occur without the other: When touching an object without moving it, we perceive pressure, but not friction. When letting a pen loosely glide over a piece of canvas we feel the vibration induced by the motion, but not pressure. This vibration is sufficient for us to experience the texture that the pen is gliding over. This idea is supported by the modern understanding of the physiology of tactile perception: There are four main types of nervous receptors in the skin. Ruffini's Cylinders and Merkel's Disks are related to skin deformation and pressure perception. Meissner's Corpuscles respond to vibration from ~30 Hz to ~80 Hz while Pacinian Corpuscles react to vibration from ~250 Hz to ~350 Hz. A large body of studies suggests that perception of textures is linked to vibration at frequencies sensed by the Pacinian system [13, 14, 95, 103, 208]. This suggests that we can create an experience of texture using solely vibrotactile actuation.

7.2.2 *Vibrotactile Feedback Technologies*

As vibrations are key to the experience of texture, we review technologies for generating vibrotactile feedback. The most common way of doing so is using Eccentric Rotating Mass vibration motors (ERMs). Rumble packs for game controllers were early uses of ERMs [136]. ERMs fit into a mobile device but are typically limited to alerting, shaking, and pulsating. In research they appear to be the go-to solution for quick experimentation (e.g., [156]), though the limitations created by the slow speed up times and a coupling of intensity and frequency of the ERM stimulation are well understood [211]: the amplitude and frequency of their actuation cannot be controlled independently.

Piezo actuators have been used to overcome this limitation. While piezo elements are often used for friction reduction in haptic interfaces (e.g.,[5, 199]), they can also provide traditional vibration at lower frequencies as well as clicking sensations. Various methods have been suggested for using this to augment displays of mobile devices with additional haptic cues [105, 141, 142]. While piezo elements have high temporal precision, they have relatively small actuation range, and therefore low achievable amplitude.

Vibrotactile feedback can also be created using solenoid-style actuators (also known as voice-coils, tactors, or haptuators). Such actuators work as audio-speakers do: A magnetic core is constrained within a copper coil. The magnet moves proportionally to the amplitude and direction of the electrical signal applied to the coil. Using audio speakers for haptic feedback was first described in 1926 to enable deaf people to 'feel' speech [58]. Since then, devices have been improved to minimize sound generation [205]. These devices can be controlled with an audio signal, achieve a higher velocity than piezo actuators, and achieve high temporal precision. Therefore, they have become a popular tool for exploring haptic feedback within the HCI community, for example in papers by Israr and Zhao [211–213], Strohmeier [171] and others [64, 79, 206].

7.2.3 *Coupling User Action and Vibrotactile Feedback*

As argued above, texture is experienced through movement. Therefore, there has been a growing interest in coupling movements and vibrotactile feedback to create experiences of roughness, compliance, and other dimensions of haptic experiences. Nara et al. [123] demonstrated a 'slider' consisting of steel balls on a variable friction surface. Using a friction reduction approach, Nara et al. were able to provide distinct haptic sensations by adjusting the frequency at which they provided bursts of low friction relative to the motion of the user's finger.

Tactile texture discrimination in robotic applications is typically achieved by moving a probe over a surface and analyzing the frequency and spectral response of the signal [204]. This approach of measuring textures with a moving probe was adopted by Romano and Kuchenbecker who coupled such a recording device to a playback device. The playback device is held by the users and, as it is moved over a flat and smooth surface, provides them with the sensation of moving the device over one of the pre-recorded materials [40, 146]. This link between user action and haptic feedback need not be limited to motion. Kildal [90] explored coupling pulse speeds to pressure exerted on a surface, providing users an experience of compliance. Yao and Hayward [206] coupled pulse speed to the angle at which a rod is tilted, providing an experience of an internal rolling stone. Strohmeier et al. [171] presented a flexible device which couples pulse frequency to changes in the amount by which the device is bent, resulting in an experience of changing material composition. All work listed above is based upon a common principle: When coupling vibrotactile feedback with user motion, vibration and motion are perceptually combined, leading to a new experience. The vibration is no longer attributed to a vibrating actuator, but rather is felt to be a property of a dynamic system that does not vibrate [123]. Therefore, if one wishes to find parameters of vibrotactile feedback that lead to an experience of texture, these parameters must also be adapted to user motion.

7.2.4 *Open Questions*

The literature suggests that coupling vibrotactile feedback with user motion is promising. However, previous work has either not systematically varied the parameters with which the feedback is generated [171, 206], or presented only anecdotal results regarding the mapping of feedback parameters to experiences of texture [123]. Kildal conducted a qualitative study of two levels of four parameters (granularity, amplitude, grain-distribution and regularity), demonstrating that they could create a variety of sensations [90]. However his analysis was not designed to link variations in feedback parameters to variations in experience of texture. Kildal stipulates that "Future controlled studies will focus on answering this question." [90], p.7. We next describe a simple haptic interface that we use to conduct such an experiment.

7.3 IMPLEMENTATION OF HAPTIC FEEDBACK DEVICE

We envision vibrotactile feedback coupled to human action to be used for augmenting tangible interfaces with additional dynamic material properties – similarly to how projection is used to augment the appearance of tangible tokens. Ideally we would like to explore such feedback in unconstrained space, using 3D motion tracking. However, for the sake of a controlled experiment and to maximize spatial and temporal

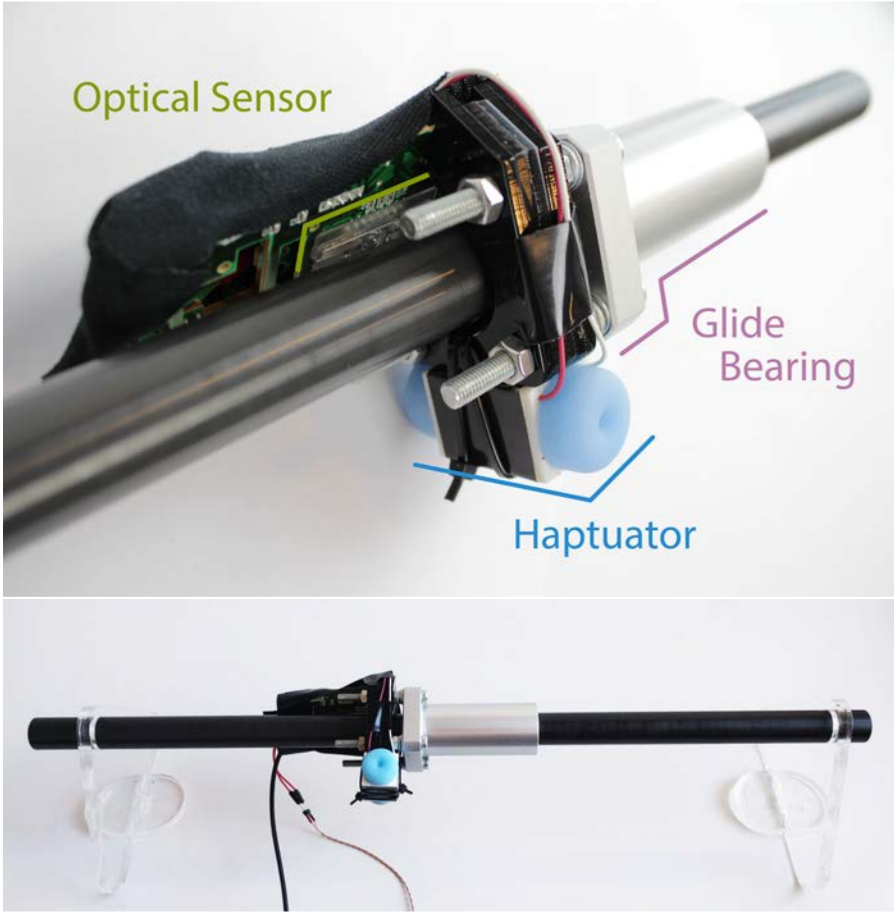

Figure 24: Slider used for experiment. Participants interact with slider by moving the silver glide bearing. The optical sensor measures the movement and the Haptuator vibrates the device relative to the speed at which the slider is moved.

sensing resolution, we constrain interaction to moving an object along a straight path.

MECHANICS We created a custom slider using a linear glide bearing (length 80 mm) with a Frelon GOLD® lining and a anodized aluminum rod (⌀ 20 mm, length 500 mm) as seen in Figure 24. We opted for glide bearings, because they create only negligible vibrations when moved. Sensing

We used an optical sensor, harvested from a Logitech M500 mouse. The sensor was placed on the gliding element, so it would move a computer's cursor as the slider was moved. The update rate of the sensor was measured to be 125 Hz and the step resolution was measured to be 0.032 mm (note that we disabled mouse acceleration).

SIGNAL FLOW & VIBROTACTILE ACTUATION We used a BM3C Haptuator by TactileLabs. The haptic feedback was generated as an audio signal based on cursor position, which was sampled at 200 Hz using

Max/MSP. The audio signal was played back using an external soundcard (UR44 by Steinberg) and a low power, generic audio amplifier connected to the haptic actuator. The delay between onset of motion and haptic actuation was estimated to be around 20ms (4.5ms from the soundcard, 2ms for registering movement, 8ms for updating cursor position from the mouse, and 5ms from Max/MSP operating in overdrive mode).

PARAMETERS OF VIBROTACTILE ACTUATION Sound is typically produced by a vibrating object that causes longitudinal waves in the air, which we then hear. Because what we hear is directly linked to the vibration of such an object, we can use existing vocabulary that describes sound for describing the vibrations that cause the sound, such as amplitude (loudness), frequency (pitch), and timbre (the quality of a sound, or color – consider the difference in sound between the vowels in 'eek' and 'oh').

While amplitude and timbre can directly be applied to haptic feedback, frequency cannot be a feedback parameter, as we vary frequency with user input speed. Instead we use granularity as a constant that, multiplied with the motion of the user, results in frequency. We generate haptic feedback as a series of 64 sample pulses. The frequency with which they occur is based on the user's action and granularity of the virtual texture:

$$f = Granularity \times UserMotion$$

User motion is defined as the speed with which the slider is moved. It is measured in cm per second:

$$UserMotion = \frac{cm}{seconds}$$

granularity represents the number of features on a surface. We defined it as pulses per cm (p/cm):

$$Granularity = \frac{pulses}{cm}$$

The pulse frequency can therefore be expressed both by granularity multiplied by user motion, or for the implementation, as pulses per second:

$$f = \frac{pulses}{cm} \cdot \frac{cm}{seconds} = \frac{pulses}{seconds}$$

Finally, we pass the signal through a bandpass filter to modulate its timbre, allowing us to create sensations which are qualitatively distinct while sharing the same granularity and amplitude.

7.4 EXPERIMENTS

Having established the feedback parameters, we are now interested to investigate how these can be used to create different texture experiences. There are a number of established psychophysical research methods that are used to "derive an understanding of the relation between changes in the physical stimulus and the associated sensation" [85], p.11. Of those, we chose to use magnitude estimation, in which users estimate the strength of individual stimuli by assigning numbers to them [60, 162]. Because this method does not set a predefined maximum or minimum, we felt that it was best suited for an experiment in which the presence of the target experience is not known. The result of this experiment will allow us to create response curves that show how a change in the vibrotactile feedback influences the experience of texture.

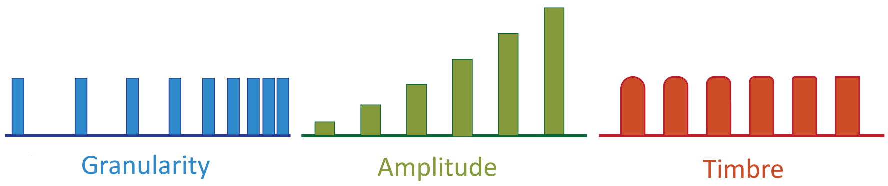

Figure 25: Naïve visualization of vibrotactile parameters.

To make basic comparisons between the effects of individual parameters and to validate the magnitude estimation experiment we decided to add a second task. Participants were asked to produce the texture experience that they previously evaluated, using the same parameters as in the magnitude estimation task.

As we were interested in potential interaction effects of the haptic feedback parameters, we opted for a factorial design. Based on a pilot study, we chose to compare the effects of 5 levels of granularity, 3 levels of amplitude and 4 levels of timbre on participants' experience of roughness, bumpiness, adhesiveness and sharpness.

7.4.1 *Experimental Apparatus*

We used the linear slider described in the implementation section. The experimental flow and data-logging were done in Processing. Communication between Processing and Max/MSP was handled by OSC [201].

To ensure that participants base their responses solely on their haptic experience, participants were asked to wear headphones during both tasks of the experiment. The headphones were playing white noise to mask any external sound.

7.4.2 *Independent Variables*

When reasoning about the effects of different feedback parameters, and for choosing appropriate levels, we think of them as shown in a naïve model in Figure 25.

GRANULARITY We imagined granularity to correspond to individual surface features. When impulses can be distinguished from each other, we expected them to be described as bumps. At higher granularity levels, for which individual pulses cannot be clearly distinguished, we expected users to report an experience of roughness. We expected to find a point at which pulses are generated so rapidly that they cannot be distinguished from each other at all, leading to smooth experiences, potentially influencing the perceived adhesiveness.

The granularity levels we chose were 312.5, 19.53, 4.88, 2.44, 1.22 pulses per cm. The choice was constrained by the sensing resolution of the optical flow sensor used (0.032 mm per step). Our particular software implementation also required us to use values sharing a common denominator. Based on the sensing resolution, the highest achievable granularity was 312.5. The other values were chosen based on a geometric series, while still having a distribution that naively felt equidistant to the experimenters. We chose to pick geometric series as these reflect our acoustic understanding of frequencies: octaves form a geometric series (e.g., A3 = 220 Hz, A4 = 440 Hz, A5 = 880 Hz).

AMPLITUDE We expected amplitude to modify the intensity of a given experience, while not having any influence on the type of experience. Amplitude levels chosen were set in Max/MSP to -9.8db, -6.8db and -3.8db relative to line level. If the amplitude approached line level any closer, there were some timbre and granularity combinations which could make the experimental apparatus vibrate to the extent that the optical flow sensor could detect the vibration. This would lead to a feedback loop causing continuous vibration. The lowest value was chosen so that all combinations would still be clearly perceivable. The medium value was selected halfway between these two (in regards to sound, the perceived amplitude doubles every 6db).

The output from Max/MSP was set to default, as was the internal volume regulation of the UR44. The output of the UR44 was set to 75 % and connected to a 4.5V, 1W preamp set to maximum volume.

TIMBRE We believed that timbre would have an influence on how clearly impulses can be felt, interacting with how granularity is experienced. We also expected timbre with a high frequency peak to feel sharper than timbre that peaks at low frequencies. We adjusted the timbre of the pulse train using a band-pass filter. The filter was implemented using the state variable filter object (svf∼) of Max/MSP with

Figure 26: Objects used to discuss texture experiences.

the Q set to default. We chose to center the filter on 40, 80, 160 and 320 Hz (in the rest of the paper, when we speak of 'high' or 'low' timbre, we are referring to the center frequency of this filter). These values were chosen as they encompass both the typical response frequencies of Meissner's Corpuscles (~30 to ~80 Hz) as well as the Pacinian system (~250 to 350 Hz) and because they are a geometric series.

7.4.3 *Experimental Measures*

The dependent variables were the participant's estimation of roughness, bumpiness, sharpness and adhesiveness. To ensure that we had a shared understanding of the words chosen to describe these experiences, we discussed them using example objects (Figure 26). We described adhesion as a measure of stickiness which is highest when the slider felt most sticky and is lowest when the slider did not feel sticky or felt slippery. We demonstrated this by the difference felt when moving a finger over the smooth area of a stone, compared to the silicone surface of a bicycle light. We intend it to capture the sticky/slippery dimension described by Okamoto [130].

We described roughness as a sensation relating to how coarse a texture is. Roughness is lowest when structures are very close together, as if the slider was moving over very fine sandpaper and higher when structures are larger and further apart, as if moving over coarse sandpaper. We discussed this using the broken edge of the stone in Figure 26 and the smooth side of the stone as examples. Roughness is used to capture the micro-roughness dimension [130].

We described bumpiness as the experience that there are distinct shape features on the object, which could be distinguished from others. Low bumpiness is when the slider feels as if it is moving over a flat surface, high bumpiness is when there are a large number of shape features. We again used the stone which had several bumpy features as an example. The comb seen in Figure 26 was used to discuss that as bumps move closer together, they might no longer be experienced as discrete bumps. Bumpiness is expected to capture macro-roughness [130].

We described sharpness relating to bumps as an estimate of how pointy a bump is. For experiences of roughness, we described sharpness as 'the potential of the texture to scratch you'. We used the pointy and blunt sides of the comb as well as sandpaper and canvas as examples. Sharpness was added based on feedback from participants in a pilot study.

7.4.4 *Task 1: Magnitude Estimation*

In this task we investigate how the perception of textures changes when the stimulus changes (e.g., "Does roughness increase with granularity?" or "How does changing the timbre influence how adhesive something is perceived to be?"). To do so we conducted a fully factorial magnitude estimation experiment based on the design suggested by Stevens [162] as described by Gescheider [61]. For each trial we varied levels of frequency, amplitude or timbre. The measures were the user's estimation of adhesiveness, roughness bumpiness or sharpness. The measures were blocked and the blocks were counterbalanced between participants. The feedback parameters were randomized for each block.

Participants were read the following text (adapted from Gescheider [61]) and given a written copy, which the experimenter discussed with them sentence by sentence.

"As you move this slider, we will provide you with varying haptic stimuli. Your task is to tell us how strongly you experience (adhesion/roughness/bumpiness/sharpness) by assigning a number to the sensation. Call the first sensation any number that seems appropriate to you. Then assign successive numbers in such a way that they reflect your subjective impression. There is no limit to the range of numbers you may use. You may use whole numbers, decimals or fraction. Try to make each number match the intensity with which you perceive the sensation."

Participants were told not to set a maximum or minimum value before they started the experiment and were instructed to report their initial judgements without dwelling too long on any particular trial.

Figure 27: Participants conducting the magnitude estimation task.

7.4.5 *Task 2: Haptic Texture Production*

This task investigates how the perception of textures compare to each other (e.g., "Is a high frequency timbre component more important for the experience of sharpness than for the experience of roughness?" or "Does amplitude play an equal role for all experiences investigated?").

Participants were presented with a digital interface with 3 sliders: A 5-point slider for frequency, a 3-point slider for amplitude and a 4-point slider for timbre, corresponding to the levels of the stimuli experienced in task one. The sliders were not labelled and the participants received no instructions on the effect of moving the sliders. Participants were asked to create the sensation they felt best represented roughness, bumpiness, sharpness or adhesiveness to them. They were not given a time limit. The discrete scales were chosen, so participants would not be able to create experiences of texture which they were not presented with during the magnitude estimation task.

7.4.6 *Experimental Procedure*

Upon signing of consent forms the experimenter discussed the dependent variables with participants as outlined above. Once participants and experimenter felt they had a shared understanding of the experimental measures, the experimental procedure was explained. When the participants felt that they understood the instructions, they conducted a practice experiment with 7 combinations of levels for each measure. This was done to familiarize participants with the device and prevent learning effects.

Task one and two were interwoven. After participants completed a block of task one, they proceeded to create the corresponding experience for task two. Each participant spent approximately 75 minutes on the entire procedure (Figure 27).

7.4.7 *Participants*

We recruited 24 participants of which 10 were female. Participants were between 22 and 79 years old (M 35.8, SD 13.6).

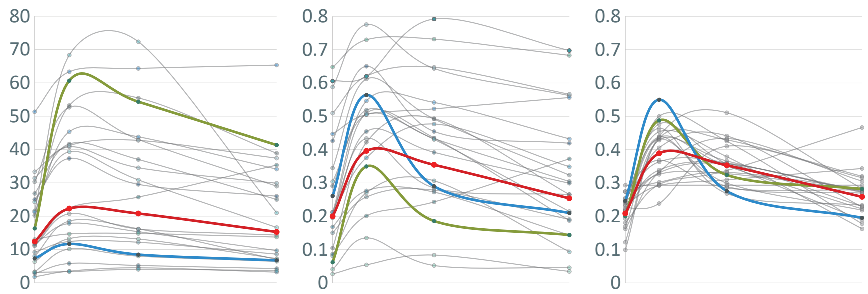

Figure 28: Data processing steps: means of raw data (left), normalized geometric means (center) and visual response scales (right). Each line represents the data of one participant. The geometric mean of all participants is indicated in red.Two participants are highlighted to show how the transformations influence individual response curves.

7.4.8 *Data Analysis*

The raw measures (Figure 28, left) of task one (magnitude estimation) were normalized per participant, by dividing each participant's data by their highest response (Figure 28, center), as discussed by Jones et al. [85]. A repeated measures multivariate ANOVA was conducted on this normalized data. For descriptive statistics we took the geometric mean for each level of each parameter, as suggested by Gescheider [61], p.239. For creating visualizations of the data, we translated the individual response scales, based on the difference of participant average from the grand mean average (Figure 28, right), as suggested by Han et al. [68]. For the exact calculation please refer to the spreadsheet provided with the supplementary material, or see Han et al. [68]. Note that we cannot make any claims in terms of magnitude between the responses of individual participants or between the strength of the experiences. What the response curves do show is if and how a change in our haptic feedback parameter (x-axis) led to a change in how participants experienced the texture (y-axis).

The data for task two required no further normalization. We analyzed it using a within-subjects multivariate ANOVA.

All reported statistics use Greenhouse-Geisser correction if the assumption of sphericity is violated. If the Greenhouse-Geisser estimate of sphericity is > 0.75, Huynh-Feldt correction was used. Post-hoc tests were Bonferroni corrected.

7.5 RESULTS – MAGNITUDE ESTIMATION

Based on the multivariate ANOVA, we found that manipulating haptic feedback parameters did indeed lead to changes in the experience of texture. We found significant main effects for timbre ($F_{12,204} = 8.100$, $p < .001$) amplitude ($F_{8,88} = 6.647$, $p < .001$) and granularity ($F_{16,368} = 2.942$, $p < .001$). The way the experience of texture is influenced, differs for each feedback parameter: granularity does not exhibit a clear trend, which is reflected in its low effect size ($\eta_p^2 = .113$) while timbre had a larger effect ($\eta_p^2 = .323$) that appeared quadratic. Amplitude appeared linear and had the strongest effect size ($\eta_p^2 = .377$). We did not find any interaction effects. To better understand the experiences of texture, we next look at the individual univariate results. Because our main focus is to better understand the experiences of texture, we will report the rest of the results grouped by experience type, as visualized in Figure 29.

7.5.1 *Bumpiness*

While granularity did not have a significant effect on bumpiness (Figure 29, top), we can see a negative trend. Bumpiness is strongest for granularities below 4.88 p/cm. The effect of granularity appears nonlinear. It drops between 2.44 p/cm and 4.88 p/cm but otherwise the negative trend appears negligible. We consider bumpiness to be equivalent to macro-roughness, and as such expected bumpiness to increase where roughness decreases. The response curve of granularity indeed shows an opposite trend to roughness (Pearson's r = - 0.84). Amplitude had a significant effect on perceived bumpiness ($F_{1.185,27.479} = 41.567$, $p < .001$) though the effect size was somewhat lower than for roughness ($\eta_p^2 = .644$). The response curve for timbre peaks at 80 Hz and then directly starts to decline. The effect of timbre was significant ($F_{2.235,51.410} = 25.006$, $p < .001$, $\eta_p^2 = .521$) and had a quadratic response curve. Timbre at 80 Hz and 160 Hz was different from 40 Hz and 320 Hz ($p < .005$).

7.5.2 *Roughness*

In Figure 29 (second from top) we see that perceived roughness increased with increasing granularity ($F_{1.905,43.819} = 6.170$, $p < .005$). While overall this effect is not particularly strong ($\eta_p^2 = .212$), for the lower range of granularity it was experienced much stronger than for higher levels. Below 4.88 p/cm the average rating increase per p/cm is 3.46% of the grand mean, while above 4.88 p/cm it only increases by

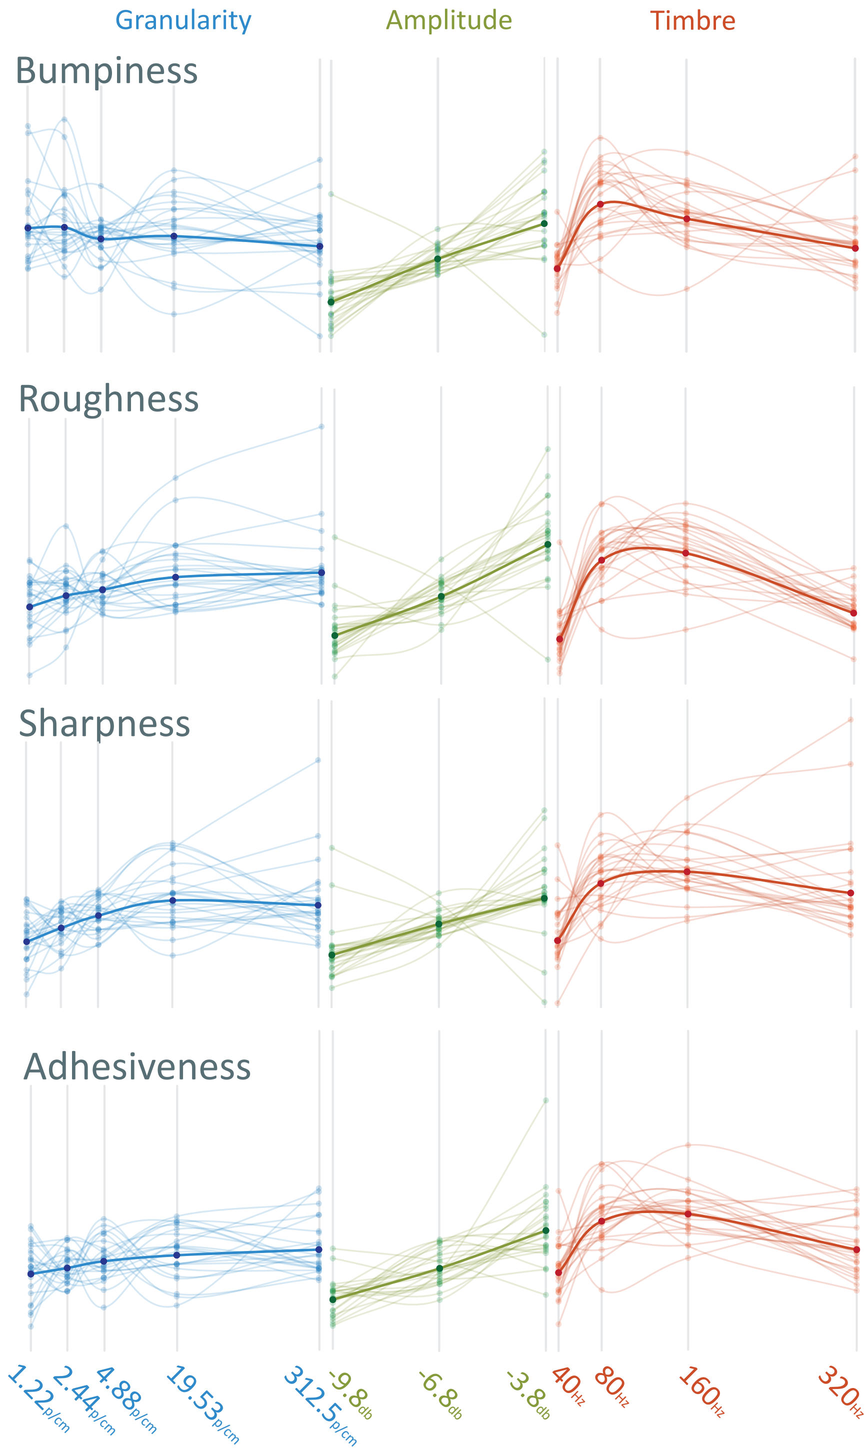

Figure 29: Response curves. Each line represents the geometric mean of how strongly a participant experienced roughness, bumpiness, sharpness or adhesiveness at the indicated level of the feedback parameter. The bold lines represents the geometric mean of all participants.

0.04%. This effect is somewhat hidden by the logarithmic scale: please note that the granularity values double with each step on the x-axis.

We expected that as granularity increased, it would eventually be experienced as smooth and that we would see a dip in our response curve (blue, left). While this did indeed happen for some participants, the mean actually slightly increased. We also did not find a significant difference between a granularity of 19.53 p/cm and 312.5 p/cm. This could mean that we did not test a wide enough range of granularities, or that the timbre levels that were experienced as rough masked the effect of the decreased granularity. Looking at amplitude, we can see that there is a strong linear effect on perceived roughness ($F_{1.216,27.976} = 49.357$, $p < .001$, $\eta_p^2 = .682$). Post hoc analysis revealed that all levels were significantly different from each other ($p < .005$).

Finally we can see that timbre has a steep rising slope between 40 Hz and 80 Hz. While for 8 participants the experience of roughness peaked at 80 Hz, the grand mean continued to rise by an additional 5.45%, peaking at 160 Hz, after which the experience of roughness declines. Timbre had a significant effect ($F_{2.554,58.743} = 49.063$, $p < .001$, $\eta_p^2 = .681$) and post hoc analysis revealed that 40 Hz and 320 Hz were different from 80 Hz and 160 Hz ($p < .005$). A contrast confirmed the quadratic nature of the response curve.

7.5.3 *Sharpness*

Sharpness (Figure 29, second from bottom) appears superficially similar to roughness, but there are slight differences which will become more prominent in task two. Low granularity was typically not experienced as sharp; there appears to be a significant positive trend ($F_{2.49,57.259} = 7.913$, $p < .001$, $\eta_p^2 = .256$). Beyond 19.53 p/cm this effect is weaker, though for 11 participants the experience of sharpness continued to increase at 312.5 p/cm. Amplitude had a significant effect on sharpness ($F_{1.110,25,522} = 17.86$, $p < .001$) but the effect size is much lower than for roughness and bumpiness ($\eta_p^2 = .426$). The response curves for timbre show that the experience of sharpness declined relatively little between 160 Hz and 320 Hz (for sharpness the decline is 18% of the grand mean, while for adhesiveness it is 24, for bumpiness it is 32% and for roughness it is 46%). In fact, for 9 participants the experience of sharpness increased between these two values. This suggests that high frequency timbre is most likely to lead to sharp sensations. The overall effect of timbre on sharpness was also significant ($F_{2.164,49.772} = 11.903$, $p < .001$, $\eta_p^2 = .341$). Sharpness again lead to a quadratic response curve, however, it was the only experience for which the timbre level of 320 Hz was not significantly different from 80 Hz and 160 Hz.

7.5.4 *Adhesiveness*

Looking at all response curves for adhesiveness, there appears to be high agreement. However, many participants reported that they had difficulty rating adhesiveness. Because of this, we suspect that the lower amount of variance simply means that there were very few moments at which users felt adhesiveness strongly enough to give it a confidently high rating (Figure 29, bottom). We found statistically significant effects of granularity ($F_{4,92} = 4.770$, $p < .005$, $\eta_p^2 = .172$), amplitude ($F_{2,46} = 32.212$, $p < .001$, $\eta_p^2 = .583$) and timbre ($F_{3,69} = 15.841$, $p < .001$, $\eta_p^2 = .408$).

7.6 RESULTS – TEXTURE PRODUCTION

In task two the activity of the participants is inversed. The experiences of texture, which so far have been the dependent measures, have now become a single independent variable, and the feedback parameters which previously were independent variables now become the dependent measures. As expected, we found that experience type had a significant effect on how participants used granularity ($F_{3,69} = 2.87$, $p < .05$), amplitude ($F_{2.15,61.808} = 7.618$, $p < .005$) and timbre ($F_{3,69} = 4.892$, $p < .05$). Figure 30 shows the number of participants that chose a particular level of a parameter. Note that the peak for high amplitude may appear exaggerated compared to other scales as participants have fewer options to select from. In general, the results from task two agree with the results of task one.

7.6.1 *Bumpiness*

Bumpiness (Figure 30, top) shows that most participants favored the low granularity levels, though surprisingly 6 participants chose 312.5 p/cm. Amplitude again confirms the previously observed effect, and for timbre most people also chose 80 Hz, as expected (compare to Figure 29, top).

7.6.2 *Roughness*

As expected for roughness (Figure 30, second from top), few participants chose low granularity levels. The high number of people who chose 312.4 p/cm is also in agreement with our results from task one, but contradicts what we expected. Amplitude confirms the previously observed strong effect, while for timbre most people chose 160 Hz (compare to Figure 30, second from top).

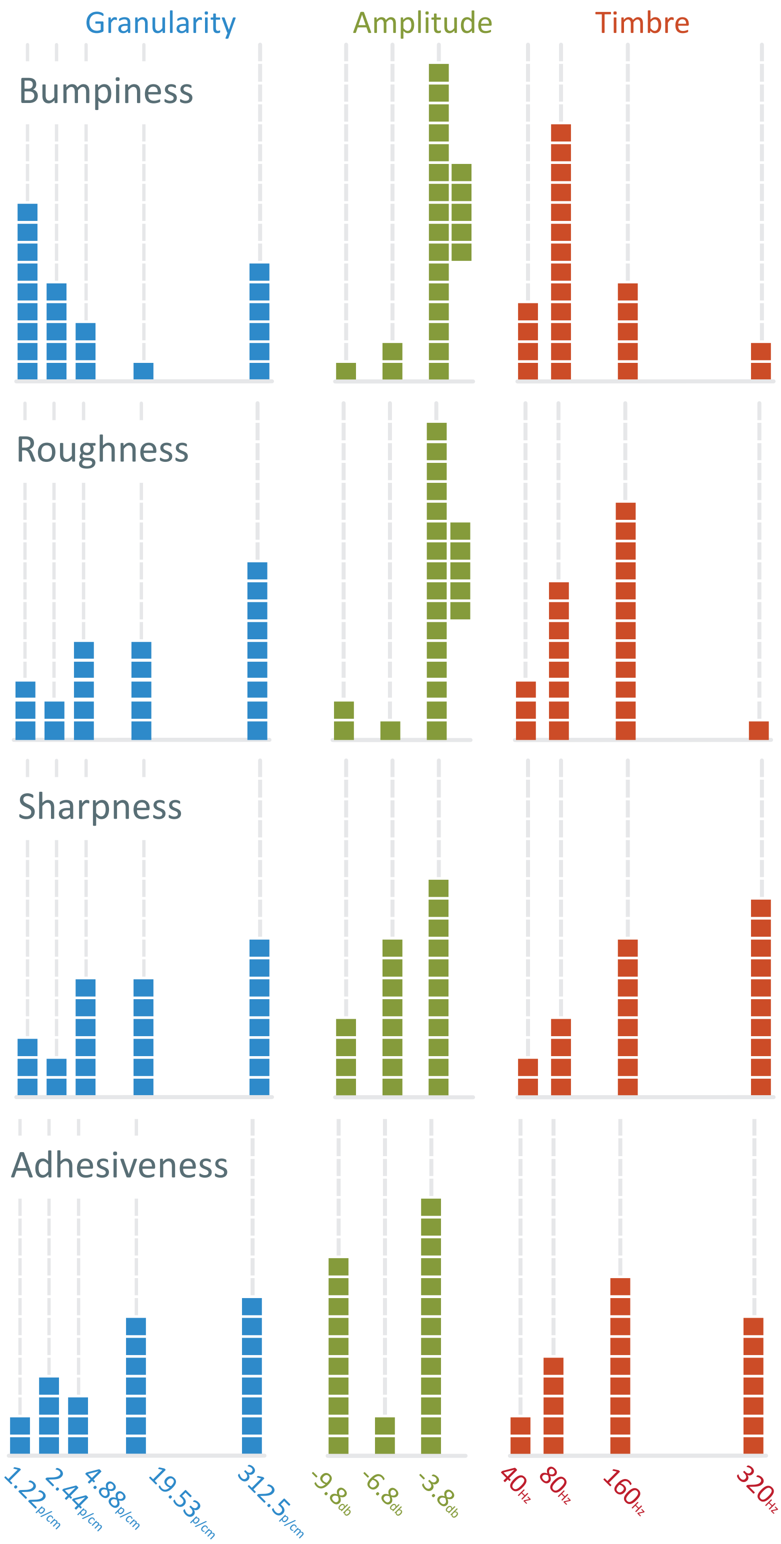

Figure 30: Texture production results: each square represents a participants' choice of a parameter when generating a the corresponding haptic experience.

7.6.3 *Sharpness*

For Sharpness (Figure 30, second from bottom) more people preferred high granularity. We clearly see that amplitude is less important, as participants distributed their choice comparatively evenly. Most participants chose the highest timbre value, which was expected based on task one (compare to Figure 30, bottom left).

7.6.4 *Adhesiveness*

The stand out feature for Adhesion (Figure 30, bottom) is that participants had split opinions on amplitude. Almost half felt that lower amplitude lead to a stronger experience of adhesion. granularity and timbre were used somewhat as expected, with participants trending towards higher levels (compare to Figure 30, bottom).

7.7 DISCUSSION

The aim of this study was to investigate how actuation parameters relate to haptic experiences. In particular, we have focused on the experiences that can be created from coupling actuation parameters to movement. Next we discuss the main findings on this coupling.

7.7.1 *Actuation parameters and haptic experiences*

We found that amplitude had a very strong effect, but that it was not equally important for all experiences of texture. Contrary to our expectations, we found a much weaker effect of granularity; however, it does appear to play an important role in distinguishing between micro and macro textures. Finally, we introduce the concept of timbre, which has received very little attention in the context of vibrotactile feedback so far. We found that it had a relatively strong effect and that, within the constraints of the sampling points we collected, it had a quadratic response curve. Our data also indicates some interactions between timbre and granularity which were not intuitively obvious to us. We found that the 312.5 p/cm level of granularity had surprisingly high levels for roughness and bumpiness in both tasks. We believe that this was caused by timbre overriding the effect of granularity: while neither roughness nor bumpiness had a granularity level which they were uniquely correlated with, bumpiness was clearly associated with a timbre of 80 Hz and roughness was clearly associated with 160 Hz. We believe that participants who optimized for timbre in the texture production task chose the highest granularity level as this maximizes the effect of timbre. Conversely, when participants experienced texture with high frequencies, the effect of timbre overrode the effect of granularity.

timbre can also be used to stimulate a participant at a fixed frequency without the experience of vibration: while an object that receives pulses at 160 Hz is experienced as vibrating, an object that receives pulses relative to its motion ($f = Granularity \cdot UserMotion$) while resonating at 160 Hz is not perceived as vibrating. This may provide interesting opportunities for future studies on haptic perception.

The strong effect that timbre had also suggests that there is value in haptic impulses that do not map linearly to user motion. Recent research on haptic perception also suggests that our haptic experience is not linearly related to how fast we move relative to an object [24]. We expect future work to explore alternatives to the linear mappings that have been used so far.

7.7.2 *Methodology and Limitations*

Like other magnitude estimation studies on haptic perception [12, 13, 81, 93], the number of levels of independent variables greatly affect the results that can be obtained. As we were interested in capturing interaction effects, we further constrained our number of levels by choosing a factorial design. As we do not anticipate any of the mappings between feedback parameter and experience of texture to be linear, this provides a clear limitation. The precision of our results for granularity and especially for timbre could have benefitted from more levels.

The combination of magnitude estimation and texture production proved interesting, despite each participant merely producing a single texture per experience. Constraining the participant's options to the same levels for both tasks allowed us to easily compare them. Using a continuous scale instead would have allowed us to find the true peaks of the different sensations, which we would like to explore in future work. Magnitude production appears particular appealing for exploring the coupling of actuation parameters and movement, because of the extent to which the sensation is produced by the participants themselves.

7.7.3 *Using the Results*

While the experiment contributes directly to an understanding of haptic experiences and action-coupled vibrotactile feedback, there are also a number of potential immediate applications of our results. For example, Valve recently released its SteamVR Tracking Hardware Development Kit (HDK) [83]. This HDK enables augmenting virtual reality experiences with custom objects and controllers, which can be augmented with feedback as we describe it. For example, in a virtual kitchen, you could feel the difference between cutting on wood and cutting on stone. While playing virtual golf, you might feel the texture of sand or grass as your golf club touches the ground.

The haptic feedback device that we use is similar to those used in high-end mobile devices [6]. When navigating a foreign city, the methods described in this paper could provide directions by subtly changing the 'feel' of directions, for example, making moving north feel smoother than moving east or west. We see this type of haptic feedback as adding to the repertoire of methods available for the design of Tangible Interfaces. From their inception Tangible Interfaces have been augmented with additional modalities, be it projection [182], shape change [143] or dynamic material properties [121]. Our exploration adds a micro-dimension of digital material surface features, with the intent to move the future of HCI one step further away from 'pictures under glass' [23].

7.8 CONCLUSION

We presented a method of generating vibrotactile feedback relative to the user's motion. We demonstrated that this method is able to convey the experience of texture when manipulating a tangible object. Our data suggests that roughness and bumpiness can be separated by granularity while sharpness and adhesiveness appear to be experienced when timbre levels are higher. Roughness is also associated with lower timbre than bumpiness, and both roughness and bumpiness are more dependent on amplitude than sharpness and adhesiveness are. This relation between haptic textures and vibrotactile feedback was demonstrated to be consistent both when participants perceived a texture and had to evaluate it, as well as when participants were asked to create a texture. The findings in this paper can be applied in applications using commodity hardware, as tracking technologies and high-end devices with the necessary haptic actuators are becoming more common.

7.9 ACKNOWLEDGEMENTS

We would like to thank Mathias Petræus for his support in conducting the experiment and Jarrod Knibbe for continuous feedback on this project.

This work was supported by the European Research Council, grant no. 648785.

8 ON MAGNITUDE ESTIMATION

8.1 REVISITING MAGNITUDE ESTIMATION

Magnitude estimation – as used in our paper *Generating Haptic Textures with a Vibrotactile Actuator* [170] (Chapter 7) – is one of the most common psychophysical scaling methods [60]. It is usually associated with Stevens [162] and has been used since the 1950s for establishing the relationship between the magnitude of a physical stimulus (such as light or sound) and its perceived strength. Experimenters have also explored its use as a method in HCI research. For example, McGee [116] demonstrated the use of magnitude estimation for usability ratings. While initially conceived to study the relation of a basic stimulus to its perceived magnitude, magnitude estimation has since been used repeatedly to explore both complex stimuli with multiple parameters [170] and their relation to complex concepts such as danger [49] or uncertainty [9].

If little is known of how the data relates to the magnitude of interest, most common scaling methods used in HCI become problematic. For example, using a binary scale, one assumes that there is a meaningful distinction between stimuli. If there is ambiguity, it will not be captured. Likert-style items would be used based on the assumption that the granularity chosen is appropriate for the stimuli of interest; however, it is possible that all stimuli may be perceived to be in a single category – meaningful differences might still be present, but not captured. Using a semantic differential scale, one makes the implicit assumption that the two poles are diametrically opposed, which might not be the case. Furthermore, using any bounded scale with a maximum and minimum assumes that there is a known maximum and minimum which the participant can clearly identify. Otherwise one might encounter a situation where the participant experiences a stimulus they wish to rate stronger than all previous stimuli, but the scale does not allow the selection of a higher item. This results in bunching of the values at the high or low end of the scale.

If one uses magnitude estimation instead, one makes very few assumptions on the relation between the stimulus and the target estimate and can be confident that the result will indeed reflect the magnitude of interest, rather than artifacts created by the data collection method: Magnitude estimation captures ambiguity and allows participants to chose the granularity with which they report effects freely. Magnitude estimation makes less implicit assumptions on the nature of the relation between stimulus and experience than a semantic differential scale would. Magnitude estimation also avoids clustering at the extremes of a scale, as participants can always extend their scale. Most importantly, magnitude estimation can capture both small and large effects simultaneously.

In personal, informal, communication with authors who have used magnitude estimation, the typical reason for using it is that magnitude estimation is "quick" or "easy".

Other reasons for choosing magnitude estimation are the assumption that the results are on a ratio scale [109, 116], that it is simple to implement [157], that it can capture suprathreshold effects which would not be captured by discrimination tasks [181], and, finally, that it can adapt to the individual preferences of the participant: the participant is free to use a binary scale, to use a 10-point scale, or to simply assign values without stipulating an underlying scale [145].

8.2 LITERATURE REVIEW

To better understand how magnitude estimation is currently being used within the HCI community, we conducted a literature review in the ACM Digital Library. While we originally intended to constrain our search to papers published at CHI, we extended our search to the entire Digital Library, due to the low number of such studies.

A thorough analysis of the opportunities and limits of magnitude estimation as an experimental method in general was also presented by Gescheider [60] in 1988.

Searching for the term "Magnitude Estimation" resulted in 22 hits. After reviewing the abstracts, we removed four papers from the sample as they were technical evaluations. And two other papers as they did not actually use magnitude estimation. One hit referred to an abstract without a full text version. Two papers were removed as they re-used data collected in an experiment presented in a third paper. Finally, two pairs of paper were near identical in content, so the duplicates were removed.

The resulting set consisted of eleven papers published between 2004 and 2018. Of these eleven, six presented the results of the magnitude estimation as their core contribution. These included:

- finding polarity and magnitudes of sonification parameters for audible data representations in general [191], as well as mappings for *error*, *uncertainty* [9], *stress* and *danger* [48];
- mapping parameters of vibrotactile feedback to texture experiences [170];
- finding polarity and magnitude of vibration parameters for tacton design [49]; and
- magnitude of perceived image motion during head movements [50].

Five papers used the results of the magnitude estimation study for further domain-specific purposes. These included:

- ranking the usability of systems [116, 150];

- finding features that impact the readability of sentences [157];
- creating a model of V1 of the occipital cortex [181]; and
- assessing methods for using crowdsourcing for creating relevance scores [109].

The papers investigating usability and crowdsourcing presented the use and analysis of magnitude estimation itself as a contribution towards their respective fields.

Strohmeier and Hornbæk present the results as response curves, providing insight on the shape of the relation between stimulus and estimate [170]. Other papers report only the slope of a linear regression. The papers by Walker [9, 191] and Fergusson [48, 49] further report the results stratified by observed polarity.

Most papers implement the experiment as suggested by Gescheider [61]. Often even the experiment instructions follow the same or similar wording [48, 49, 109, 116, 150, 157, 170]. The exceptions are papers that omit exact descriptions [9, 191] or intentionally change the task to suit their needs better [50, 181].

Sometimes experimenters provide participants with a reference stimulus. They might say "This stimulus has a value of 100, please rate the other stimuli relative to that value". Gescheider discusses the pro's and con's of this [60].

However, in detail the papers diverge on methodology. Some papers use one [150, 181], two [116] or even three [50] reference stimuli during the experiment. Typically participants are asked to use any number larger than zero. Others explicitly allow zero as a response [157], while others provide no bounds at all [49, 150]. Most papers then either log-transform the data or use the geometric mean. Two use the arithmetic mean [157, 181], and one paper uses the median [50].

Of those papers that report on prepossessing (8/11), two do not preprocess the data [48, 181], two papers map the data to a range from 0 to 1 [48, 170], two express the results as a percentage of either a fixed value [50] or the participant's mean response [109], and two papers removed the task or participants' means from the estimates [116, 170].

Of those papers that conduct an ANOVA (5/11), this is twice done on the raw data [48, 191], once done on the rescaled data [170], once on the aligned rank transformations (See [200]) of the rescaled data [49], and once on z-scores [157].

A full comparison of methods and other relevant details can be found in figure 31.

8.2.1 *Problems with magnitude estimation*

While there are many benefits to using magnitude estimation within HCI, the method is not without its problems, and it is not universally applicable. For example, when the goal is to briefly assess the usability of a system, Sauro et al. politely suggest that the added complexity introduced by the method, compared to a Likert scale or the like, is probably simply not worth it. [150].

Title	Method	Anchors, Range	Measure of Central Tendency	Preprocessing	Statistical Tests	Purpose	Result Presentation
Master Usability Scaling: Magnitude Estimation and Master Scaling Applied to Usability Measurement (2004)	Standard	Two anchors, unbounded	Geometric Mean	Subtracting of global mean	none	Further domain-specific analysis (Ranking)	n/a
Evaluation of a Multiscale Color Model for Visual Difference Prediction (2006)	Assigning a numerical value to the difference between images	Anchor Stimulus set to 20	Mean	none	none	Further domain-specific analysis	n/a
Perception of Image Motion During Head Movement (2006)	Custom	Three anchor stimuli, (used to linearize the results)	Median	Rescaled to percentage of maximum anchor, analyzed on medians	none	Magnitude Estimation	Slope, Descriptive Statistics
Comparison of Three One-question, Post-task Usability Questionnaires (2009)	?	Anchor stimulus, unbounded	?	?	none	Investigating use of ME	n/a
Universal Design of Auditory Graphs: A Comparison of Sonification Mappings for Visually Impaired and Sighted Listeners (2010)	Standard	>0	Geometric Mean	?	ANOVA on raw data, geometric mean for slopes	Magnitude Estimation	Slope
Displaying Error & Uncertainty in Auditory Graphs (2012)	?	?	Geometric Mean	?	Correlation	Magnitude Estimation	Slope
Offline Sentence Processing Measures for Testing Readability with Users (2012)	Standard	>=0	Mean, Z-Scores	Standardization	ANOVA, Post Hoc on zScores	Further domain-specific analysis (Ranking)	n/a
Generating Haptic Textures with a Vibrotactile Actuator (2017)	Standard	>0	Geometric Mean	Mapped from zero to one, means removed	ANOVA on normalized data, mean removed for slopes only	Magnitude Estimation	Curve
On Crowdsourcing Relevance Magnitudes for Information Retrieval Evaluation (2017)	Standard	>0	GeoMean	Data confirmed log normal, individual results expressed as a percentage of grand mean	none	Domain-specific analysis, investigating use of ME	n/a
Evaluating Mapping Designs for Conveying Data Through Tactons (2018)	Standard	Any range of numbers, as a multilplier of the original stimulus	Geometric Mean	Mapped from zero to one	ANOVA on aligned rank transforms	Magnitude Estimation	Slope
Investigating Perceptual Congruence Between Data and Display Dimensions in Sonification (2018)	Standard	>0	Geometric Mean	none	ANOVA, Posthoc Analysis	Magnitude Estimation	Slope

Figure 31: Overview of analyzed papers. "?" indicates that the information was not reported in the paper, "n/a" indicates that the question does not apply. "ME" is short for magnitude estimation. Most typical features highlighted.

Even in situations where it is clearly appropriate, there are a number of problems with how magnitude estimation is currently being used.

Problem 1: Results are difficult to interpret

Typically, magnitude estimates are presented either as summary statistics [116, 157, 170] or as slopes [9, 48, 49, 191]. When looking at simple summary statistics, it is difficult to judge if the difference between two stimuli is small or large, as one is not familiar with the scale. So while the ratio-scale nature of the collected data is often cited as the reason for choosing magnitude estimation, this benefit is often lost when presenting the data. The reader might infer that the difference between stimulus A and B is smaller or larger than the difference between B and C, but has no indication of the magnitude of these differences.

In papers which present the slope, the information the reader receives is even less meaningful. While the polarity of the slope provides a qualitative indicator which might be useful, the slope itself is close to arbitrary. Consider that average values of 2, 4, and 8 for stimulus levels 1, 2, and 3 would produce a different slope from average values of 20, 40, and 80 for the same levels. This would make the results appear different, even though the ratios are identical. Ignoring this for a moment, the reduction to a slope is even more problematic, as it suggests that effects are linear, which is usually not true in real world situations. In fact, in real world situations, where one might wish to optimize a stimulus according to various constraints, or find an optimal combination of stimuli, these non-linear effects are much more important than the general trend. Additionally, the slope one finds is strongly determined by the range of stimuli chosen. If, for example, I estimated loudness as a function of frequency, I would find a positive slope between 1 Hz and 1000 Hz and a negative slope between 1000 Hz and 100,000 Hz (consider also Anscombe's quartet, shown in the margin).

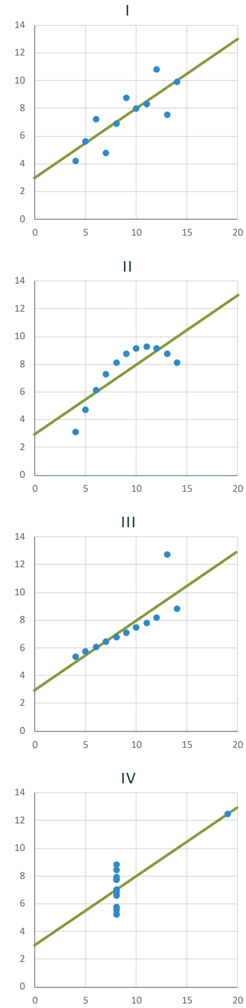

Anscombe's quartet: four datasets with near identical simple descriptive statistics. The regression line of each set is $y = 3 + x/2$. The differences between sets I,II, and III are relevant in magnitude estimation.

Problem 2: No standardized method of analysis

Even if one is satisfied by simple summary statistics or slopes, there is no consensus on how to calculate them. Several attempts have been made to formalize the analysis process. However, these are either vague [61] or do not fully consider all practical implications of running such an experiment [68]. Assuming that the estimates are truly ratio-scale data, they can be averaged while preserving the original ratio. However, the ratios estimated by the participant using the largest scale will have the strongest effect on the final data. This problem could be addressed by rescaling the data. However, depending on how this is done, the rescaled data might no longer preserve the original ratio scale. Depending on the range of values participants can chose from, and the type and order of pre-processing steps, this can lead to final results which might look similar, but require very different interpretation.

Even if two papers used the same prepossessing steps, chances are that such a comparison would result in observing differences or similarities among participants, rather than effects of the stimuli. Master scaling – providing one or more common stimuli across experiments – has been suggested as a method to address this issue, but this is often undesirable [61] or impractical, for example when comparing effects on different modalities.

Problem 3: Finding a Baseline

The original intention of conducting a magnitude estimation experiment was to measure the "sensation magnitude" [60]. Within the context of HCI, though, magnitude estimation has been used to collect mappings for haptic and visual data representations, investigate qualitative differences in experience, and assess the appropriateness of sentence structures and entire documents. Participants in such experiments are no longer simply asked how strongly they experience a stimulus, but rather how strongly the experienced stimulus matches some criteria. This requires additional considerations in the data analysis process.

The response curves we found (see Chapter 7, Figure 29) were very similar for all experiences. During my CHI talk, it was suggested that this might be because the haptic actuator has a higher amplitude at specific frequencies or that the experimental device resonates more strongly at some frequencies. While we cannot rule this out, I believe the similarities we found are due to the receptive properties of the Pacinian corpuscles.

For example, if the experimenter wishes to design haptic icons to present a number of data variables [49], they might conduct a series of magnitude estimation studies, establishing a mapping between a set of haptic parameters and each data variable. To implement the results, though, it is less important to know what the absolute mapping between haptic parameters and each data variable is, than to know which haptic parameters allow participants to best differentiate between data variables. In such a situation, one might treat the mean results over all data variables as baseline, and conduct the analysis on how the haptic parameters for each data variable differ from that mean.

If an experimenter wishes to find the ideal audio frequency for communicating danger, they will, naturally, find that frequencies between 20 Hz and 20.000 Hz will be rated higher than frequencies outside of that range. Also experimenters will most likely find high ratings between 2000 Hz and 5000 Hz where hearing is most sensitive [59]. However, if one is interested in designing a "danger" sound, it is probably obvious that one would want this sound to be perceived as loud as possible. An interesting question might however be "ignoring how loudly one experiences a sound, which frequency is perceived as most indicative of danger?"

While further experimental work is needed to confirm this, it seems reasonable to assume that participants use both the salience and the quality of a stimulus in a magnitude estimation task. If only the magnitude of experience is being studied, it's not necessary to differentiate between these two concepts. However, in the more complex tasks found in an HCI context, it might be beneficial to consider if one is interested in salience, quality, or both, and adjust the study design accordingly.

8.3 REVISITING THE ANALYSIS PROCESS

While there are obvious drawbacks to using magnitude estimation, the minimal assumptions made about the relation between stimulus and estimate still make it an appealing experimental method, especially in exploratory HCI studies, where the parameter space is not well understood. Taking full advantage of such studies requires a reassessment of the analysis process in the context of the more complex stimuli and target experiences which might be used in an HCI study.

The following section provides a detailed discussion of what such an analysis might look like. Using the dataset collected in *Generating Haptic Textures* [170] and the wisdom of hindsight, I demonstrate a step-by-step analysis of complex magnitude estimation data which avoids the problems I just outlined.

8.3.1 *Dataset*

The dataset collected for the *Haptic Textures* experiment (see Chapter 7) consists of 5760 entries. 24 participants[1] each provided 240 estimates in total. Participants could individually create their own scale 4 times, once for each target experience (Roughness, Bumpiness, Sharpness and Adhesiveness). For each experience, we counterbalanced 3 levels of amplitude (low, medium, high), 4 levels of timbre (40 Hz, 80 Hz, 160 Hz and 320 Hz) and 5 levels of granularity (1.22, 2.44, 4.88, 19.52 and 312.23 pulses per cm). No anchor stimuli were used.

The dataset is available in the ACM Digital Library ***https://dl.acm.org/citation.cfm?id=3025812.***

The final dataset consists of 5760 estimates between 0 and 200 (M = 22.36, SD = 22.75), of which 9 estimates are invalid. Please refer to section 7.4 for further details on the experiment and data collection.

8.3.2 *Sources of Variability*

The assumption in a magnitude estimation task is that the stimulus being manipulated relates to the experience in a set and repeatedly measurable way. While we assume that there is such a clear relation, the data does not obviously reveal it without carefully considering all sources of variability and removing those irrelevant to the research question at hand.

USER UNCERTAINTY Users might be unsure of what they are experiencing and consequently have variability in their assessment, or they might be unsure of how to asses the experience, even if they feel sure about what they are feeling, or both. This type of uncertainty would lead to higher variability in the ratings. If this uncertainty is higher

1 In the re-analysis, we removed participant 24, as their ratings are in strong disagreement with the other 23 participants, and evidence emerged that they may not have understood the task.

than our signal of interest – the relation between experience and stimulus – we are unable to demonstrate that such a relation exists. The lower the variability is in relation to our signal of interest, the higher the confidence we can describe it with. While the ratings provided to us by participants are unit-free, the ratio between variability based on stimulus and variability based on user uncertainty provides us with a tool to describe how strong the change in experience is. I will from now on refer to user uncertainty as noise. It is important to note that *I make the assumption that this noise is relatively similar for all users.*

The assumption of similar noise need not strictly be met. If there is a user who can produce estimates with less noise, this merely results in larger confidence interval. That person's estimates also has a proportionally stronger impact on the results.

DIFFERENCES BETWEEN USERS Users might also asses stimulus differently in terms of magnitude. Differences in physiology and experience might make a stimulus more or less salient – salience being the strength with which a stimulus is perceived. For example, the age or gender of the participant can influence the absolute magnitude of judgements [60]. This might result in a constant difference between users. Similarly, even if users experience the stimulus with the same salience, they might still rate them using a difference reference point or scale, which additionally adds a relative difference between users. This type of variability between users is irrelevant for us and we can remove it.

COMMON EFFECTS The conditions of the experiment itself might influence the ratings in a systematic way. In our own case, the frequency of the vibrotactile stimulation is fundamentally intertwined with its salience. The same is true for the frequency of sound. Consider the example I used before, where one wishes to find out what frequency to use for an audible danger notification. If we conduct a magnitude estimation experiment where participants assess how strongly a sound is indicative of danger, will the result merely reflect how loudly the participants perceived the sound, or would we find a frequency which somehow has an intrinsic quality that communicates danger? Ideally, such factors should be considered in the experimental design, for example by recording a baseline and reporting on the difference to the baseline. Alternatively they require special consideration during the analysis process.

In our case, it appears that either timbre interacted with amplitude or that timbre interacted with the perceived strength of the stimulus. As we are interested in the differences between textures, rather than the overall strength of the experience, such effects will be removed in our re-analysis of the data.

8.4 ANALYSIS STEPS

Based on these considerations, I will provide a detailed walk-through of how a magnitude estimation experiment might be analyzed, using the data from the *Haptic Textures* experiment.

8.4.1 *Data Preprocessing*

As each user created their own scale, and because each user could chose different scales for different target experiences, we cannot meaningfully compare raw estimates from different experiences or different participants. (An example of raw data from participant 21 can be seen in Figure 32, top). Preprocessing is required before all data can be combined.

One approach is to normalized the data, so that all results are mapped to a fixed range [49, 170] as seen in Figure 32, middle. This leads to a series of problems. (1) If a user does not experience changes, they might still have small variations in their estimates. These small variations are inflated to the same range as the estimates of a user who experiences strong, clear changes. (2) If a user produces a clear scale which reflects their experience in great detail, but has one large outlier, the rest of their scale will be compressed into a small range. (3) The variability of the data is in the same units as the data. Though all user estimates are within the same range, they still are not on the same scale. Thus, the variability of the data is not homogeneous, which violates the assumptions of many of the most common statistical analysis methods. (4) Ratios are not preserved.

An alternative to normalization is standardization (Figure 32, bottom). Standardization means subtracting the mean from a dataset and dividing by its standard deviation, for example as done by Siddharthan [157]. Standardization has several benefits over normalization: it is more robust to outliers, its data retains its ratios, and it does not inflate small variations in signal. However, in the standardization process, we do not differentiate between variability due to noise or variability due to changes in stimulus. This inflates the ratings of a user who does not report any changes in stimulus, and deflates rating by a user who does experience changes. So while the overall variability is constant, the makeup of this variability varies across users and experiences. User estimates therefore still do not share a common scale, as a constant

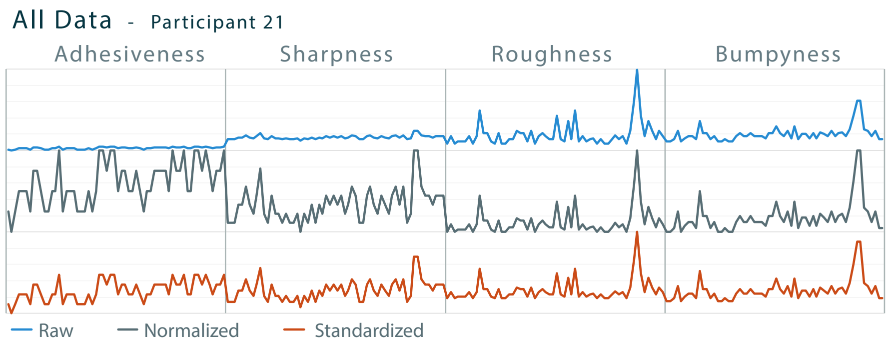

Figure 32: Raw data (top, blue), normalized data (middle, gray), standardized (bottom, orange)

difference between two estimates for two users might not represent the same perceived change in stimulus relative to noise.

The data collected per participant has a known structure, and we can use this to standardize the data relative to noise, while ignoring any variability due to changes in stimulus. This allows us to scale user estimates relative to the users' ability to consistently evaluate an experience. In this new scale, a fixed difference in estimates between two users also represents a fixed change in stimulus relative noise.

The first two steps of this procedure are similar to the calculation of the mean sum of squares for error (MS_{error}) for univariate ANOVA or the mean sum of squares for within-group variability (MS_w) for repeated measures ANOVA. The main difference is that we perform the process for each participant and each experience independently.

Based on these considerations, we perform the following preprocessing steps:

1. *Remove variation due to stimulus:* Average the estimates for each level of stimulus and subtract this average from the estimates.

2. *Calculate standard deviation:* Calculate the mean sum of squares and divide it by the appropriate degrees of freedom. This provides us with a standardized measure of noise.

3. *Standardize according to noise:* Divide the raw estimates by the standardized noise measure. The resulting estimates have a fixed signal-to-noise ratio for all participants and all experiences.

This standardization method not only provides us with a method of making comparisons between users and experiences, but also provides the basis for what we can later use to create meaningful units in our plots. Currently, all estimates share a signal-to-noise ratio. We can calculate how strong a change in signal needs to be for us to say that it is statistically significantly different from noise. When comparing stimuli, this unit would be the minimal difference the estimates need to have for us to confidently say that this difference in estimate resulted from a difference in experience.

One could also calculate the confidence interval of the noise and use half the confidence interval's range as the divisor. The resulting data might then be interpreted in the following way: Differences >1 are due to changes in experience, differences <1 might merely be noise.

8.4.2 *Finding a measure of central tendency per participant*

To analyze the data, we must calculate an estimate per level and participant. The task as described by Gescheider [61], and the analysis suggestions by Han et al. [68], assume that participants will rate values on a ratio scale. Accordingly, Gescheider and Han et al. suggest using the geometric mean as a measure of central tendency.

While we followed the previous literature in setting up our experiment, we had the impression that our participants struggled with using geometric scales. For example, if users assigned values to two stimuli, for example, 20 and 80, and encountered a third stimulus which they perceived to be halfway between, they might place the value equidistant to the two previous, creating a linear sequence (20, 50, 80) rather than a geometric one (20, 40, 80). We also noted that users struggled with a geometric series when values became very small or very large, and would instead tend towards linearly increasing or decreasing scores.

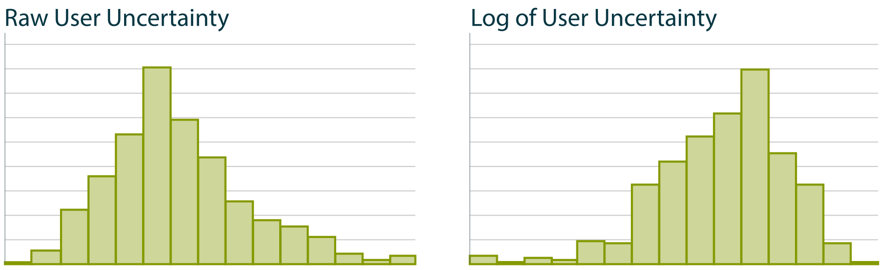

Figure 33: Histograms of user uncertainty. Left: based on raw data. Right: based on log-transformed data

This observation left us uncertain if the geometric mean, as suggested by the literature, or if the arithmetic mean, as per our intuition, is the right choice. To make a decision based on data, rather than on intuition, we analyze the per participant variability, removing any effects of stimulus or target experience. We did this analysis twice, once on the raw data and then again on log-transformed data. We use the log-transformed data, as it allows us to conduct the same calculations steps as the regular set, while preserving the geometric relations within the data. Comparing the results will allow making an informed choice for the measure of central tendency.

The steps of the analysis were:

1. *Isolation of variance due to user uncertainty:* For each combination of participant, stimulus level, and target experience, subtract the corresponding mean from all estimates. The remaining variability is what we wish to summarise with a measure of central tendency.

2. *Standardize by target experience:* For each participant, divide all estimates by the sample standard deviation for the target experience. All remaining variability is now remapped to a common scale.

3. *Find the distribution of values:* Plot the values as histograms; calculate the skewness and kurtosis of the resulting distribution (See figure 33).

If we find that the non-transformed data has a symmetrical distribution we would assume that the arithmetic mean is the preferred measure. If, however, the transformed distribution is symmetrical, we would use the geometric mean. We can calculate which distribution is more symmetrical by calculating its skewness.

$$Skew = \frac{n}{(n-1)*(n-2)} \sum \left(\frac{x_j - \overline{x}}{a} \right)^3$$

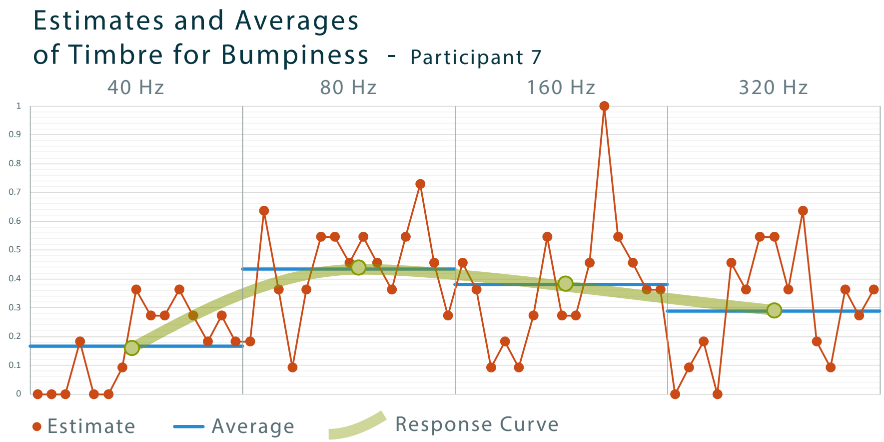

Figure 34: Data scaled so that the minimum is zero and the max is one

To further ensure that these findings were not due to outliers seen on the low end of the log-transformed data, we removed these and redid the analysis. The result did not change the overall preference for using the arithmetic mean.

We found that the raw data was indeed skewed towards positive values (skew = 0.6), which might suggest that using the geometric mean would be appropriate. However, the transformed histogram was skewed even more strongly towards negative values (skew = -1.135; see also Figure 33). While the kurtosis of the distribution would not affect our choice, it is worth noting that the kurtosis of the non-transformed dataset is closer to normal (kurtosis = 0.92) than is the transformed dataset (kurtosis = 3.94).

To ensure that this difference was not driven by outliers, we also compared each of the 96 participant-plus-target stimulus combinations between the two sets. Using the absolute values, we found that on average, the skew was lower for the non-transformed (average absolute skew = 0.17) than for the log-transformed data (average absolute skew = 0.33). We also compared the skewness case by case and found that the non-transformed data was more symmetrical in 57 of the 96 participant-plus-target combinations.

One could also decide on the measure of central tendency on a case-by-case basis, looking at each participant's data individually

We therefore performed the rest of the analysis using the arithmetic mean.

8.4.3 *Calculating per participant response curves*

Now that we have established our measure of central tendency, we can calculate how users rate the differences in stimulus. Figure 34 shows estimates provided by participant 7 for bumpiness. We show the raw data in orange. These contain the variability of amplitude, granularity, and timbre. In Figure 34, we sort the estimates according to timbre, and calculate the average for each level of timbre (blue lines). The response curve is calculated by connecting these averages (indicated in green).

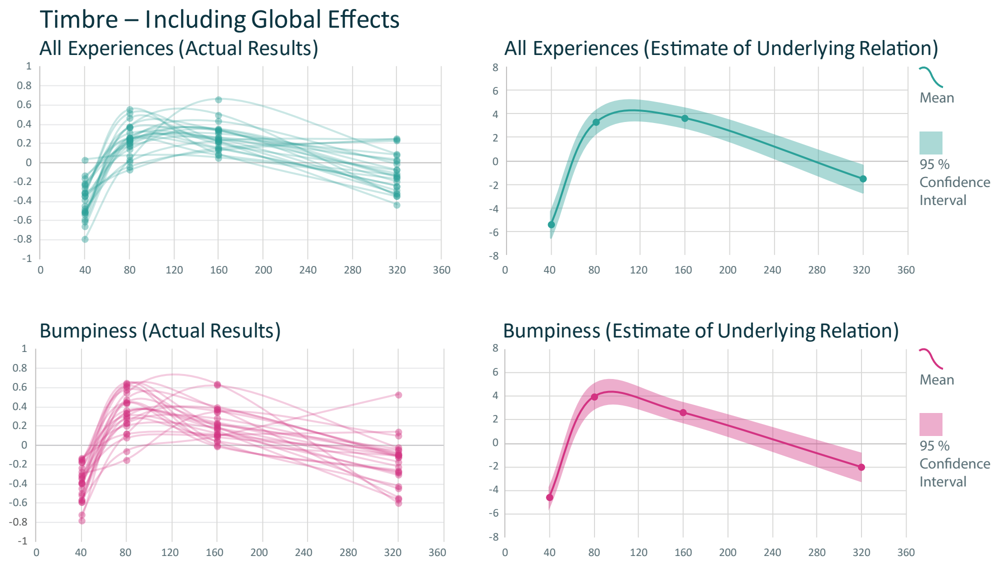

Figure 35: Without removing global effects, we can see that the results for bumpiness differ from the overall averages, however the similarity of the curves make it difficult to asses what aspects of the stimulus make an experience feel more bumpy than other experiences.

8.4.4 *Estimating underlying response curve over all participants*

We can now calculate and plot the response curves of all users to get an impression of what the overall estimates look like. As stated in the introduction, the assumption behind magnitude estimation is that the users' estimates are noisy measures of an underlying relationship between stimulus and experience. We consequently wish to estimate this underlying curve based on all users' data (see Figure 35, left). Based on the Central Limit Theorem, we know that the estimates of each stimulus level have a normal distribution, and because of the preprocessing, we also know that they have a common scale. We can therefore estimate the underlying relation between stimulus and experience by simply using the arithmetic mean. Based on our sample, we can also construct a 95% confidence interval. This range of values helps us eyeball how strong our effects are, and indicates the range in which we expect the true relation to fall, assuming our result is not significantly different from the underlying relation between stimulus and estimate (see Figure 35, right).

8.4.5 *Removing unwanted variability*

The estimated curve and its confidence interval are easier to interpret than simply plotting all user responses, as can be seen in Figure 35. The confidence interval suggests that there are clear significant differences within the ratings of bumpiness (Figure Figure 35, bottom). If we compare the response for *bumpiness* to the global response (average of responses for all four experiences, Figure 35, top) we see that while bumpiness is slightly different than the other experiences, the main sources of variability which appear to make the ratings of the timbre levels different, are due to something else than the experience of bumpiness.

In fact, the overall curve looks similar to the response curve one might expect based on the response characteristics of the Pacinian corpuscles [149, 188] (See also Chapter 3). In other words, the bulk of the variability appears to come from how strong the participants experienced the stimulus to be, and only a small portion of the variability appears to come from the way the variations in the stimuli uniquely affected their experience of bumpiness. If we wish to see how the qualities differ, we therefore need to remove the variability they have in common. We do this by comparing how strongly each target experience differs from the average response over all experiences for that level. Doing so enables us to say, "For experience X, the stimulus Y was significantly more important than for the other experiences." It will also allow us to say, "Experience A was significantly stronger for stimulus X than for stimulus Y, independently of the perceived strength of the vibration".

Again, this process may appear similar to separating variability into its components, as is done, for example, for a repeated measures ANOVA. Our intentions here, though, are quite different. When we split variability into its components with the intention of performing an f-test, we are doing so to search for the smallest valid denominator, not to demonstrate effects in the data. In fact, by removing this source of variability, we make most effects appear smaller.

To remove the unwanted variability from the data, we must calculate the marginal means per stimulus from the baseline and subtract them from the estimates of the experience we are interested in. In our case we performed the following calculations:

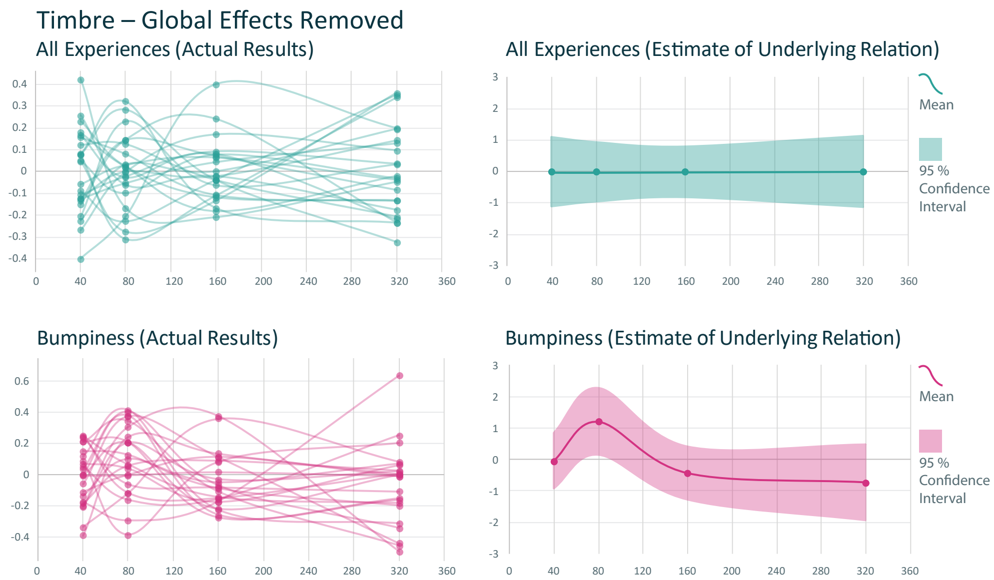

Figure 36: After removing global effects, the estimate of all experiences combined is turned into the reference that the other experiences are compared to. We now see which features of bumpiness stand out compared to the other experiences.

1. *Calculate Marginal Means:* Average the estimates for each level of each stimulus over all four experiences.

2. *Remove Marginal Means from Dataset:* For each estimate we remove the corresponding mean of amplitude, timbre and granularity.

We take these steps to make sure that any effects we might find, significant or not, are due to the source of variability we are interested in.

Figure 36 shows the same data as shown in figure 35; however, this time the global effects are removed. We see that if we average all experiences together, the resulting response curve is completely flat. The response curve of bumpy has changed between the graphs. It now shows how bumpiness is different from the other experiences. In Figure 36 we see that there is an effect of timbre at 40 Hz on the estimates of bumpiness.

8.4.6 *Rescaling to comparable Units*

As noted at the beginning of this chapter, the results of magnitude estimation studies are often difficult to interpret. While researchers present and compare slopes [48, 191] and curves [170], it's often hard to know what they mean. Typically, it is impossible to say if the difference between two slopes or points on a curve is actually meaningful, or if the difference between two values should be considered small or large.

One way of making judgements regarding the difference between two estimates is relating them to the noise that we know is present. We can then state if a difference is larger or smaller than the average noise we would expect. During our preprocessing, we re-expressed all estimates to have a fixed signal-to-noise ratio. The estimates are now all in units of standard deviation of noise.

Our ratings no longer represent a measure of the experience, but rather a measure of how many standard deviations of noise there are between the two experiences. The only thing left to do is represent this measure on a scale which is easily interpretable. To do this, we calculate the average confidence interval of our curve, and we divide our curve by that number. In the resulting curve, we can now say that if two values are further apart than 1, they are most likely significantly different from each other. This has been done in the right side graphs of Figures 35 and 36.

8.4.7 *Interpreting the Results*

We will analyze the response curve of *bumpiness* as a function of timbre (Figure 36, bottom) to explain how the graphs should be interpreted. We might first ask if Bumpiness was estimated differently than the other experiences. To do so, we check if its confidence interval includes zero. If it does include zero, we might say that finding no difference is compatible with the result we found – or, more simply put, that the two values probably are not actually different. If the confidence interval does not include zero, we have found an effect. In fact, this is what happened – the 40 Hz condition had a significant positive effect on the estimates of bumpiness, even when disregarding differences in perceived strength of stimuli (see Figure 36, bottom right).

We can also look at the rest of the curve for bumpiness (Figure 36, bottom right), and eyeball if changes in the curve appear significant. The changes from 40 Hz to 80 Hz and from 80 Hz to 160 Hz do indeed appear significant, while there appears to be little difference between 160 Hz and 320 Hz.

If we want to know for sure, we can calculate the confidence interval of the difference. We do this using the mean difference of the estimates of the two stimuli and their pooled standard deviation. We divide our target p value by the number of comparisons we intend to make (typically levels-1) to avoid family-wise error. If the confidence interval calculated this way includes zero, we cannot say with confidence that they are different.

8.4.8 *Identifying null results, and detection threshold*

Another implicit assumption of the magnitude estimation which I have not yet touched upon is that the experience one wishes to estimate is present when performing the task. This means that if we do not find any significant effects of any of our stimuli on the target experience it is unclear how to interpret that. When no effects are found, this could be because the target experience is not present, or because the stimuli selected do not influence the target experience.

While magnitude estimation is the wrong tool if we are searching for detection thresholds, in certain circumstances it can still be used to provide an upper limit of where the detection threshold might be.

An example of this is provided in *Magnetips* Chapter 2. We studied how strongly the stimulation was perceived, dependent on how far away from the device the user's finger was. While we cannot with certainty declare where the detection threshold is, post-hoc analysis can provide a conservative estimate: We can assume that the detection level has not been reached as long as there is a significant difference between a stronger and a weaker signal. If this difference is significant, then, clearly, participants are able to perceive at least the stronger stimulus of the two. In the case from Chapter 2, the 56.6 mm level was the last value to be significantly different from weaker signals further out (Bonferroni corrected, $p < 0.01$ for 84.9, 87.2, 93.8 & 103.9). We can therefore claim that the detection threshold must be either equal to or beyond this stimulus level (See Figure 37 for the estimates in question and the corresponding confusion matrix from the post-hoc analysis).

If estimates do not find any effects, we still do not know if the effect is present but constant, or if it is not present. However, if we do find significant effects, we can use the same logic to argue that the target experience is indeed present.

a) Perceived Strength by Distance from Centre (in mm)

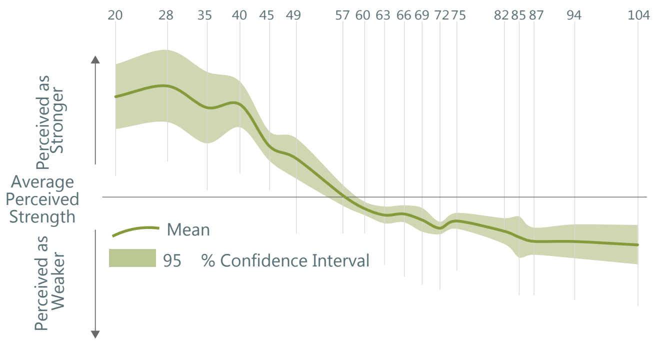

b) Confusion matrix of p-values from Bonferroni corrected post-hoc analysis of haptic feedback strength ratings by distance

Distance (mm)	20.00	28.28	34.64	40.00	44.72	48.99	56.57	60.00	63.25	66.33	69.28	72.11	74.83	82.46	84.85	87.18	93.81	103.92
20.00		1.000	1.000	1.000	.001	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
28.28	1.000		1.000	1.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
34.64	1.000	1.000		1.000	.049	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
40.00	1.000	1.000	1.000		.021	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
44.72	.001	.000	.049	.021		1.000	.001	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
48.98	.000	.000	.000	.000	1.000		.118	.001	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
56.56	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]		[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	.007	.007	.003	.001
60.00	.000	.000	.000	.000	.000	.001	1.000		1.000	1.000	1.000	1.000	1.000	1.000	.466	.462	.237	.080
63.24	.000	.000	.000	.000	.000	.000	1.000	1.000		1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	.551
66.33	.000	.000	.000	.000	.000	.000	1.000	1.000	1.000		1.000	1.000	1.000	1.000	1.000	1.000	.749	.277
69.28	.000	.000	.000	.000	.000	.000	1.000	1.000	1.000	1.000		1.000	1.000	1.000	1.000	1.000	1.000	1.000
72.11	.000	.000	.000	.000	.000	.000	.569	1.000	1.000	1.000	1.000		1.000	1.000	1.000	1.000	1.000	1.000
74.83	.000	.000	.000	.000	.000	.000	1.000	1.000	1.000	1.000	1.000	1.000		1.000	1.000	1.000	1.000	1.000
82.46	.000	.000	.000	.000	.000	.000	.125	1.000	1.000	1.000	1.000	1.000	1.000		1.000	1.000	1.000	1.000
84.85	.000	.000	.000	.000	.000	.000	.007	.466	1.000	1.000	1.000	1.000	1.000	1.000		1.000	1.000	1.000
87.17	.000	.000	.000	.000	.000	.000	.007	.462	1.000	1.000	1.000	1.000	1.000	1.000	1.000		1.000	1.000
93.80	.000	.000	.000	.000	.000	.000	.003	.237	1.000	.749	1.000	1.000	1.000	1.000	1.000	1.000		1.000
103.92	.000	.000	.000	.000	.000	.000	.001	.080	.551	.277	1.000	1.000	1.000	1.000	1.000	1.000	1.000	

Figure 37: The top figure shows us that the perceived strength declines with increasing distance. However, we have no way of inferring that the user can actually feel the stimulus over the entire range. Using the corresponding confusion matrix (bottom), we can find the range in which differences are significant – if two values are significantly different, we can assume that at least one of them was perceivable. The distance of 56.6 mm is the furthest distant which is significantly different to estimates for larger distances. It is therefore the largest distance at which we can state with certainty that the stimulus was perceived.

8.5 CONCLUSION

In this section, I have outlined many reasons why one might choose to use magnitude estimation. The most important is that in using magnitude estimation, one makes very few implicit assumptions about the concepts being explored. I have also pointed out problems with how magnitude esimtation is currently being used in HCI – most notably that there is no agreement on how to analyze the results, and therefore no clarity on how results are to be interpreted. I believe that if used with additional rigor, magnitude estimation would be an important addition to the repertoire of HCI research methods. I have thus presented a detailed evaluation of the results of a magnitude estimation task. While I perform the analysis on the data of our own experiment, I present it with the intention that the methods used might also apply to other tasks studies in HCI and hope that it will help others in their analysis. In Chapter 9, I discuss the results of this analysis and demonstrate how it helped us better understand the results of our experiment.

9 FOLLOW-UP ANALYSIS

In our paper *Generating Haptic Textures with a Vibrotactile Actuator* [170] (Chapter 7), we introduced the idea of using simple parameters for generating material experience. We described the response characteristics of these parameters, and discussed how these might be used in real world applications. We also demonstrated how participants might create four target textures using these parameters by providing the results of a magnitude production task. However, we did not explore the differences between the experience estimates, nor which parameters created such differences. Such an analysis would, however, be useful to better understand these specific experiences. To better understand the experiences of Bumpiness, Roughness, Sharpness, and Adhesiveness, I present the results of a follow-up analysis, using the method presented in Chapter 8

The data is archived with the paper and can be accessed at https://dl.acm.org/citation.cfm?id=3025812.

9.0.1 *Average Responses*

To get a better impression of how participants rated textures in general, we can look at the estimates for each level of timbre, amplitude, and granularity averaged over all experiences. Figure 38 shows these averages and the corresponding 95% confidence interval.

The unit of the y-axis is the average of the minimum distance two points need to be apart to be distinguishable at a statistically significant level for $p > .05$. Zero represents the average estimate for each stimulus.

Looking at the graph for timbre, we can deduce that the 80 Hz level of timbre was rated higher than average, and that – based on the observed variability of the data – we could clearly distinguish the observed estimate from at least three lower estimates (zero, one and two). We can also infer that, based on the observed means, we are unable to distinguish estimates of 80 Hz from estimates of 160 Hz. However, we see that the difference between 40 Hz and 80 Hz is clearly significant.

Figure 38 shows that the effects on timbre are non-linear, and the confidence interval is relatively narrow; in other words, many points are clearly significantly different from one another. However, this small confidence interval is problematic, as it suggests that individual differences between the experiences cannot be particularly large. It also suggests that some other factor is masking differences between experiences.

We selected our timbre values to fall within the responsive range of the Pacinian corpuscles (see detailed discussion in Chapter 3. We chose the 40 Hz stimulus, as we speculated that Meissner cells may play a role in the experience of Bumpiness.

Figure 38: Average Responses over all experiences for Timbre (top), Amplitude (middle) and Granularity (bottom). The x-axis represents the levels of each stimulus, the y-axis the participants estimates. Dots are actual measures, the solid line combines these mean measures and the coloured area along that line indicates the 95% confidence interval.

9.0.2 *Removing common features*

We believe that two independent factors make a participant give a high rating to an experience. The first factor, the one we are interested in, is the reaction to the question "How closely does what I feel match the target experience?" In other words, this is the participant's reaction to the quality or verisimilitude of the experience. The second factor is the reaction to the question "How strongly do I feel this sensation?" In other words, the reaction to the salience of the experience. We argue that the variability between stimuli levels seen in Figure 38 is caused by differences in the strength, the salience, with which these are perceived.

As we are interested in the quality rather than the salience of the experience, we must correct for the effects of salience. We assume that the effects common to all the experiences tested are likely due to salience, and only the relative differences between experiences can safely be assumed to be an effect of quality. Looking at the overall response curves, we see that there is a very strong effect of timbre.

In hindsight, it would have been interesting to measure the speed with which the user manipulated the experimental device, to better understand the sources of variability, and also to analyze interactions between movement and vibrotactile feedback.

This curve for timbre is similar to what we would expect to find due to the response characteristics of the Pacinian corpuscles (see chapter 3). Deviations from that curve might be explained by potential non-linear responses of the haptic actuator or the experimental device. Unsurprisingly, the curve for Amplitude shows a linear relation, suggesting that the salience increases with increased stimulation strength. The global effects of granularity are weakest. While granularity does influence the stimulation frequency, the frequency is dependant on movement. It is possible that granularity had little effect on salience, or that participants corrected for such effects by changing the speed with which they moved the slider.

Looking at these global response curves, we would assume that the response curves for the individual experiences to be very similar. In fact, we know this to be the case, due to the analysis done in the chapter 7. To truly understand what factors make a texture feel the way it feels, we must, however, focus on the variability between the estimates of the experiences, ignoring their common characteristics. We do this by subtracting the effects seen in the graphs in Figure 38.

9.1 EFFECTS OF TIMBRE ON TEXTURE EXPERIENCES

Once the variability that all experiences have in common is removed from the dataset, it is easier to see how the experiences differ from each other. Figure 39 shows how the four timbre levels we chose (Bandpass filter centred at 40 Hz, 80 Hz, 160 Hz and 320 Hz) affected the participants' estimates. The x-axis shows the frequencies we investigated. The solid line connects the average responses for each frequency. The shaded area around the solid line represents the 95% confidence interval. The units of the y-axis represent the number of distinguishable levels, given the observed noise in our data. Zero on the y-axis corresponds to the means shown in Figure 38. Positive values indicate the experience was rated higher than the overall average for that level, negative values that the experience was rated lower than the overall average. If the shaded area does not include zero, then we can state with certainty that there is a statistically significant difference of how the estimates of that experience compared to the others.

40 Hz The estimates of all experiences for the 40 Hz condition are close to zero ($<$ abs(1)). This suggests that 40 Hz did not provide any cues to participants that might help them identify one of the experiences or distinguish among them.

80 Hz Here things get interesting. Both Bumpiness and Roughness have a mean estimate of $>$ 1, while Sharpness has a mean estimate of $<$ -1. In other words, at 80 Hz, participants started to make clear distinctions between the experiences in the estimates they provided.

160 Hz Here the experiences seem to converge again. Bumpiness and Sharpness return to baseline, while the estimates of Roughness continue to be significantly higher. There appears to be absolutely no effect on the ratings of Adhesiveness.

320 Hz For the highest timbre level explored, the experiences switch and diverge in part. While increasing the timbre has no effects on Bumpiness or Adhesiveness, it had a significant positive impact on the estimates of Sharpness and a significant negative impact on the estimates of Roughness.

In summary:

- BUMPINESS estimates had a clear significant peak at 80 Hz. At 160 Hz and 320 Hz bumpiness was consistently rated lower than average.
- ROUGHNESS was most strongly experienced at 80 Hz and 160 Hz. The difference to the average rating was significant for both these frequencies, but at 160 Hz it was also the only experience which was rated significantly higher than average. Roughness estimates were below average at 40 Hz, and significantly lower than average at 320 Hz.
- SHARPNESS was rated significantly lower than the other experiences at 80 Hz and the experienced sharpness appeared to increase with increasing frequency. At 320 Hz, the estimates were significantly higher than for the other experiences.

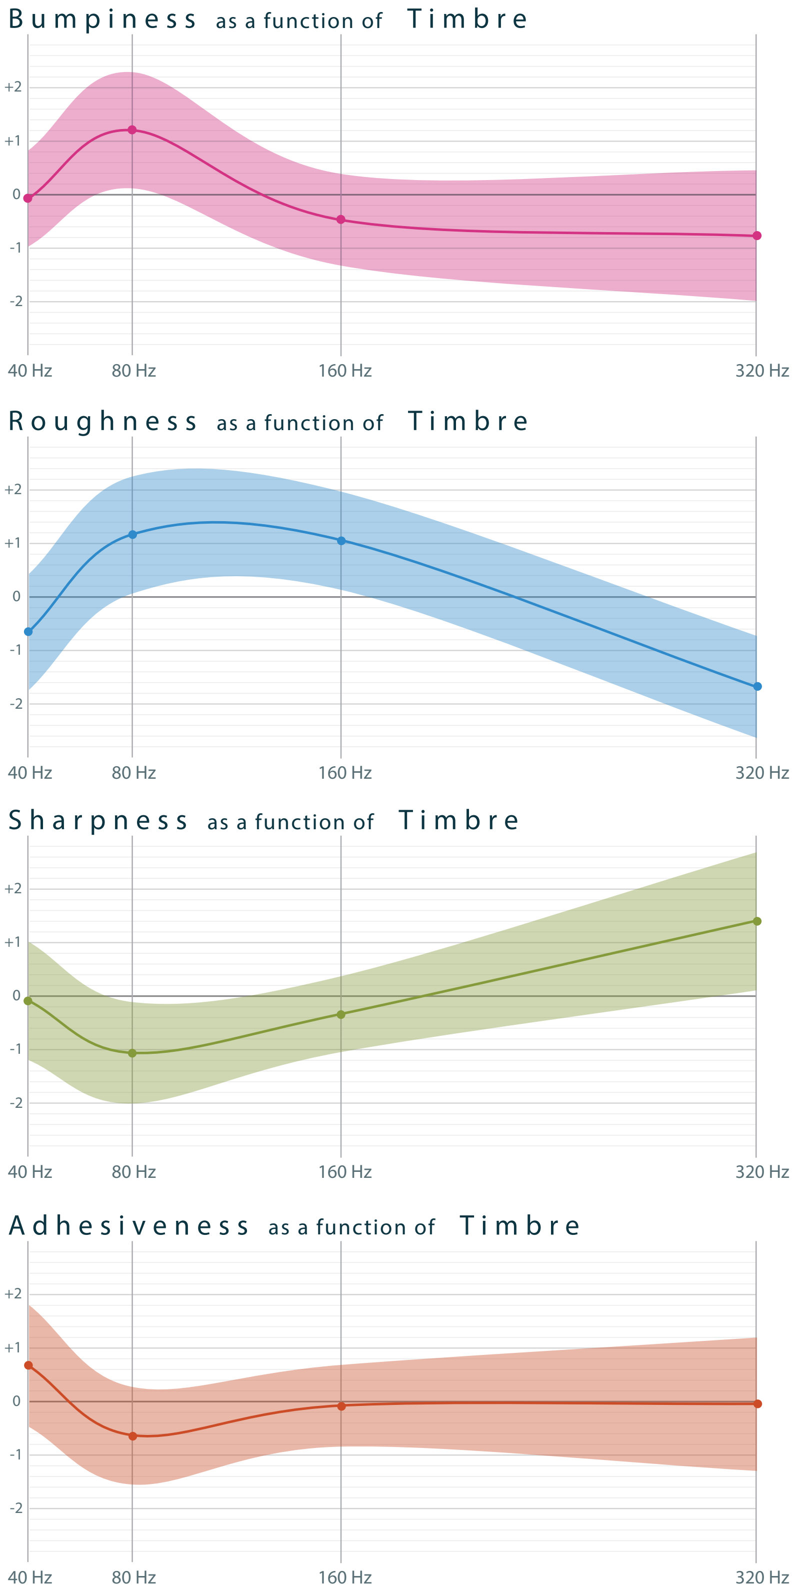

Figure 39: Average response curves (solid lines) and corresponding 95% confidence intervals (shaded area) of timbre, for each experience tested.

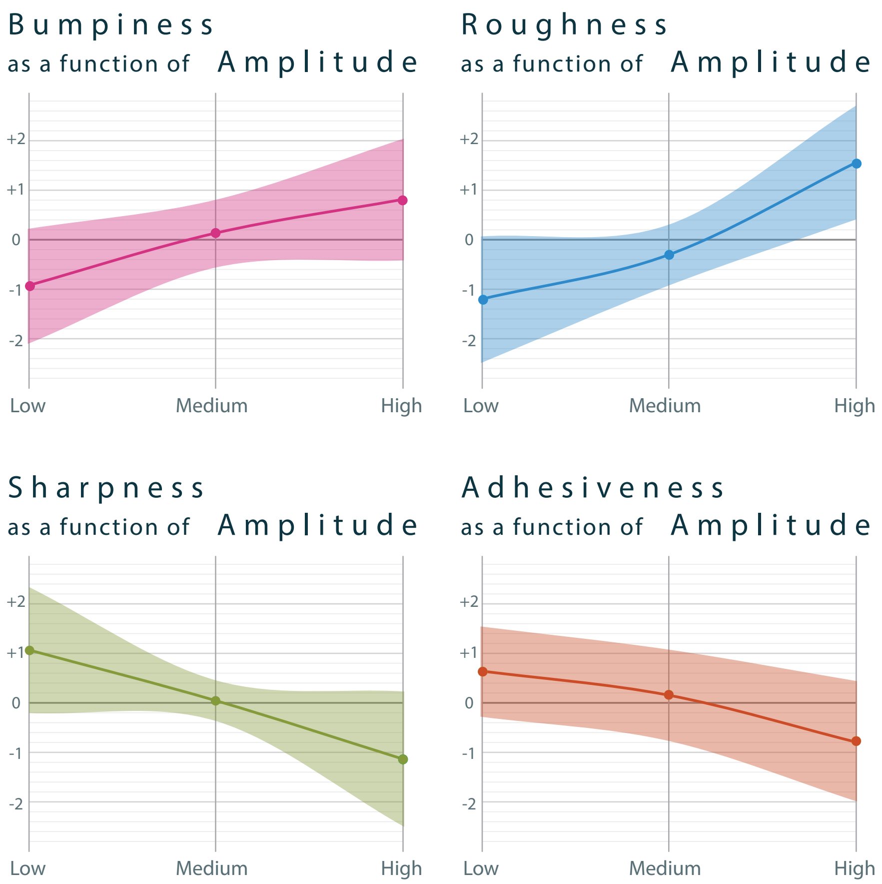

Figure 40: Average response curves (solid lines) and corresponding 95% confidence intervals (shaded area) of amplitude, for each experience tested.

- ADHESIVENESS did not lead to any significant effects on participants' estimates. Adhesiveness estimates were slightly higher for the 40 Hz condition. This effect was significant (Mean = 1.75) for low granularity (1.22 and 2.44 pulses per cm).

9.2 EFFECTS OF AMPLITUDE ON TEXTURE EXPERIENCES

We assumed that while timbre and granularity would affect the quality of the experience – whether it seemed bumpy, rough, sharp, or adhesive – Amplitude would primarily influence the salience of the experience. Looking at the ranges of the graph of responses averaged over all experiences (Figure 38), and comparing to the per-experience variability (Figure 40), we see that the effects of salience dominate those of qualia $\sim$ 2:1. However, while the effect of Amplitude is stronger than that of granularity, it is dwarfed by that of timbre ($\sim$ 3:1). There is a positive effect of amplitude on Bumpiness and Roughness and a negative effect on Sharpness and Adhesiveness compared to the grand mean. High am-

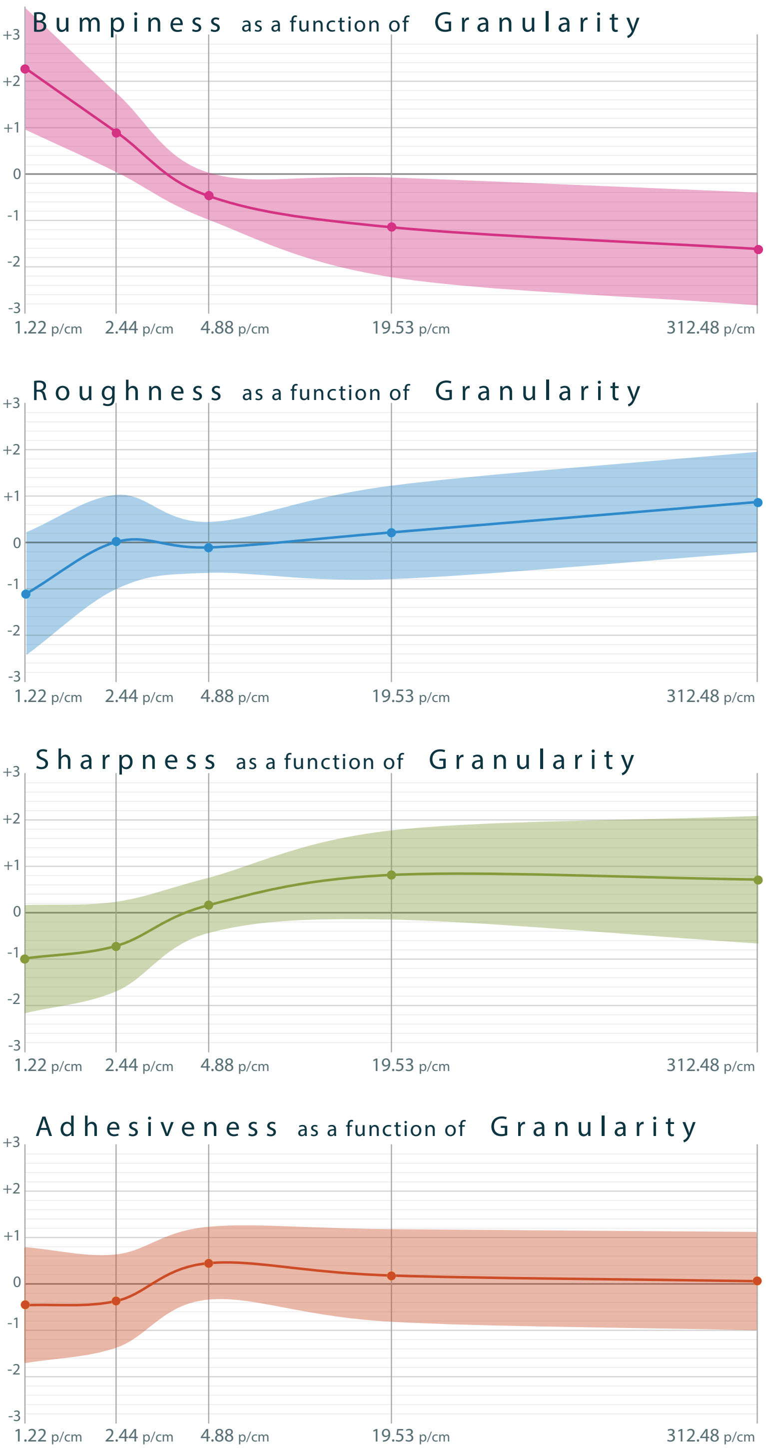

Figure 41: Average response curves (solid lines) and corresponding 95% confidence intervals (shaded area) of granularity, for each experience tested.

plitude was different from the grand mean at a statistically significant level for Roughness. This suggests that roughness not only influences salience, but also quality of an experience.

9.3 EFFECTS OF GRANULARITY ON TEXTURE EXPERIENCES

We chose the terms Roughness and Bumpiness as more colloquial equivalents for the terms Micro- and Macroroughness [130]. Based on this, we assumed that they would be clearly distinguishable on a granularity scale. Previous work by [123] also suggested that Adhesiveness might be experienced at the high-frequency end of granularity. We did not find anything of the sort.

What we did find is that fewer than 4.88 pulses per centimeter were experienced as significantly more BUMPY than the overall average, and that granularity levels above that were experienced as less bumpy. Increasing granularity levels had a positive effect on ROUGHNESS estimates, though never strong enough to distinguish it from the overall average. While SHARPNESS received lower estimates for granularities below 4.88 pulses per centimeter and higher estimates above, with a significant difference between them, these ratings did not deviate from the overall mean enough for us to identify them as distinct features of Sharpness. We did not find any effects of granularity on ADHESIVENESS.

9.4 IMPLICATIONS

We do not claim that the three parameters explored by us are sufficient, at least regularity [90] should also be considered.

This follow-up analysis confirms the results from our published paper. We can now also make new claims about our data which we were not able to do before. Because we have created a common scale based on the signal-to-noise ratio, we can make comparisons between the estimates for different texture experiences. These comparisons show that all experiences except for Adhesiveness had parameters whose ratings were significantly different than those of the others. In other words, participants demonstrated consistent differences in estimates, suggesting that the experience changed in three of the four haptic dimensions we evaluated.

Previous work has shown that material experiences can be created by recording and emulating material interactions. Our work demonstrates that a simpler, parameterized approach is also possible. This parameterized approach is not only easier to implement, it also makes working with material experiences easier in general. In creating a virtual world, the haptic experiences might be parameterized together with the visual textures. The designer will not need to hand-craft each material experience of interest, but can tweak them on a higher level. It also means that we do not need to save the information of material experiences as a raw audio file or surface profile; instead, we can simply save the descriptive parameters associated with the experience. This allows for easy sharing and on-the-fly creation of material experiences.

This follow-up analysis has also re-enforced the importance of timbre, as – even when corrected for global effects – timbre was the parameter which allowed for the clearest distinctions between textures. This points to hardware requirements for future commercial haptic devices. While Linear Resonate Actuators (LRAs) have become common place, these can only implement the granularity and amplitude dimensions. Due to their constrained resonant frequency, they cannot implement the timbre parameter, limiting their utility. Instead actuators which can provide haptic pulses over a wide frequency range are required.

10

FROM PULSE TRAINS TO COLOURING WITH VIBRATION: MOTION MAPPINGS FOR MID-AIR HAPTIC TEXTURES

Citation

Paul Strohmeier, Sebastian Boring, and Kasper Hornbæk. "From Pulse Trains to "Coloring with Vibrations": Motion Mappings for Mid-Air Haptic Textures." In:*Proceedings of the Conference on Human Factors in Computing Systems - CHI '18*. DOI: https://doi.org/32.3367/5395796.539585

Abstract

Can we experience haptic textures in mid-air? Typically, the experience of texture is caused by vibration of the fingertip as it moves over the surface of an object. This object's surface also guides the finger's movement, creating an implicit motion-to-vibration mapping. If we wish to simulate a texture in mid-air, such guidance does not exist, making the choice of motion-to-vibration mapping non-obvious. We evaluate the experience of moving a pointer with four different motion-to vibration mappings in an interview study. We found that some mappings lead to a perception-shift, transforming the experience. When this occurs, the pointer is no longer perceived as vibrating, interactions become more pleasurable, and users have an increased experience of agency and control. We discuss how to leverage this in the design of haptic interfaces.

This paper has a corresponding video-figure which I refer to in the text. It can be found here: https://youtu.be/pJ8MiFiMWHI.

10.1 INTRODUCTION

In our everyday experience, textures are always accompanied by normal force. As we move our finger over a stone wall, we push against it and the wall provides a counter-force. Research has shown, however, that many dimensions of texture experience are caused by vibration, rather than force [88]. As our fingertip moves over the stone wall, the way our fingerprint interacts with the structure of the stones causes vibration in the skin [107]. These vibrations cause Pacinian and Meissner corpuscles to fire, which in turn leads to the experience of texture [13].

It is not sufficient to simply vibrate the fingertip to make us experience the texture of a stone wall. This vibration must correlate with the motion of our finger for the material experience to emerge. We can artificially create the experience of texture, if we generate the vibrotactile feedback with a frequency proportional to the motion with which the texture is explored. This phenomenon has been investigated extensively and various haptics explorations have used this effect to manipulate the material experience of an object or create artificial textures [40, 146, 171].

Textures come with an implicit dimensionality: Moving a pencil over paper, we feel the paper's texture only when the pencil moves along its surface. If we attempt to move the pencil in a third dimension we are either constrained by the normal force when pushing down, or we no longer feel the texture when we lift the pencil. To render a haptic texture, we need to couple the vibrotactile feedback with the user's movements. If we wish to do so in mid-air, many motion-to-vibration mappings become possible. Choosing such a mapping is non-obvious.

We use a handheld device with a recoil-style haptic actuator [205] (Figure 42) to compare four motion-to-feedback mappings: no mapping (Vibration), mapping to displacement (Translation), mapping to changes in orientation (Rotation) and mapping to a point projected on a plane (Projection). The goal is to understand how motion-to-vibration mappings influence the perception of such motion-coupled vibration. We conducted an interview-based study and observed that for the Translation and Rotation conditions a perception-shift occurred: Participants described that the vibrotactile pulse trains we generated transformed into 'something more', liking the experience to 'coloring with vibrations' or 'moving through a force field'. When this shift occurred, irritation caused by vibration was reduced and simultaneously the pointing device felt as if was of 'higher quality' and moving it was 'more fun'. Users reported an increased experience of agency and a heightened sense of their body and their movements.

10.2 RELATED WORK

In this paper we discuss the experience of vibrotactile pulse trains generated by free-form movements. Our work draws on the psychophysics of touch, and is inspired by various haptic-rendering systems. In this section we discuss the physiological and technological foundation on which we build, and highlight how previous evaluations lead us to choose a qualitative, interview based approach.

10.2.1 *Texture Perception and Simulation*

The perception of texture is caused by the interaction of our fingertip with the material it is touching [107]. This interaction causes vibrations to which the Pacinian system and Meissner Corpuscles are sensitive [84]. The firing of these cells in turn is interpreted as a texture [13].

We can detect the presence of vibration within two relatively narrow frequency bands, 5 to 50 Hz (Meissner Endings) and 40 to 400 Hz

(Pacinian Endings) [84]. While this information is relatively sparse, it is sufficient for a rich set of experiences to emerge, including roughness and stickiness [13] as well as compliance [15]. Similarly to how we distinguish between the sound of two musical instruments based on the frequency profile of the tones they emit, we also distinguish textures based on the frequency profile of the vibrations cause by interacting with them [170].

This can be leveraged to create artificial material properties. For example, researchers have simulated a pen moving over a flat surface that is experienced to have the haptic properties of various other materials [40], manipulated the experienced material properties of bending an object [171], or simulated compliance for virtual buttons [90]. These simulations all used a fixed motion-to-texture coupling. We expand this work to mid-air interactions and examine the effects of various mappings on the resulting experience.

10.2.2 *Haptic Rendering Systems*

The devices used for simulating experiences such as texture or compliance typically follow two approaches. Devices using grounded haptic feedback transmit forces to the user through a kinematic chain of rigid links and joints [38]. In contrast, non-grounded feedback devices provide stimulation of the skin, but no force [39]. Alternative haptic rendering methods include body-grounded devices which provide force relative to the body, inertial approaches that transmit gyroscopic force [120, 192], or focused ultrasound [30]. Force can also be simulated by taking advantage of asymmetrical vibration [41]. A further alternative to haptic rendering is physically manipulating texture [75] or compliance [82]. In our study we use a non-grounded system, similar to the approach originally presented by Kuchenbecker [146]. This approach typically uses inertial or force sensing methods [39, 40] and detailed modelling of surfaces. In contrast, our system uses a relatively naïve model, but combines that with optical tracking, providing a large volume in which users can interact.

10.2.3 *Vibrotactile Actuators*

The ungrounded approach used in this paper requires a vibrotactile actuator. Currently ERMs are the most common solution found in products, dating back to ‘rumble packs’ used in the game-controllers of the early 90’s [136]. While ERMs are easy to implement in prototypes, other devices, such as piezo-actuators (e.g., [104]) are required for more controlled feedback. Solenoid-style actuators including Tactors, Haptuators or voice coils are typically used for texture rendering, as they can independently modulate frequency, amplitude, timbre and velocity. Using solenoid-style setups for haptic-feedback in psychology and psychophysics research is first documented in the early 20th century, using re-appropriated audio speakers [58]. Since then, research has reduced the audible and increased the tactile output of such actuators [205]. They have since found wide usage within the HCI community [171, 211–213].

These devices are usually controlled by audio-signals. Their output can be partially audible, sparking explorations of the interactions between haptic and audible feedback, for example using a handheld device that coupled audio and tactile cues based on user motion [2, 3]. As our haptic device uses audio signals for control, it exhibits many similarities with these devices. While using similar actuators to previous work [39, 40, 170], we expand upon this work by exploring new methods of designing the actuation signal.

10.2.4 *Evaluations of Haptic Experience*

Evaluating and, especially, communicating what good haptic design is, is non-trivial. This is reflected in how researchers chose to evaluate their work. The bulk of evaluations focus on detection thresholds and on studies evaluating if the device does what it is intended to do [30, 39, 40, 80, 158, 197]. From a human-centered design perspective such information provides limited value. Instead, behavioral studies that investigate how haptic feedback influences task performance are often preferred [41, 99, 142, 171, 193]. Such studies, however, do not provide a reader with insight regarding what the haptic experience feels like. In consequence, there are various studies that require participants to report on their impression of a stimulus such as object length [206], compliance [154], or roughness [170], while feedback parameters are adjusted. These studies enable a reader to understand the comparison between parameters of the specific setup, but are often not suitable for comparison to natural objects or other systems.

When research does report on the subjective experience of haptic systems, this is often done in passing [171] or as aggregated questionnaire data [25, 90]. A notable exception is an interview study by Obrist et al. [128], which presents in-depth interviews comparing haptic feedback designed to target either Meissner or Pacinian corpuscles. Because we find such research currently underrepresented, this exploration also focuses on the subjective experience. To do so we chose to use an approach inspired by Petitmengin [138, 139].

10.3 HAPTIC FEEDBACK SETUP

Typically, vibrotactile feedback is generated as a pulse train : a repetitive series of pulses, separated in time by a fixed interval (for example, the Oculus touch controller can currently produce either a 160 Hz or

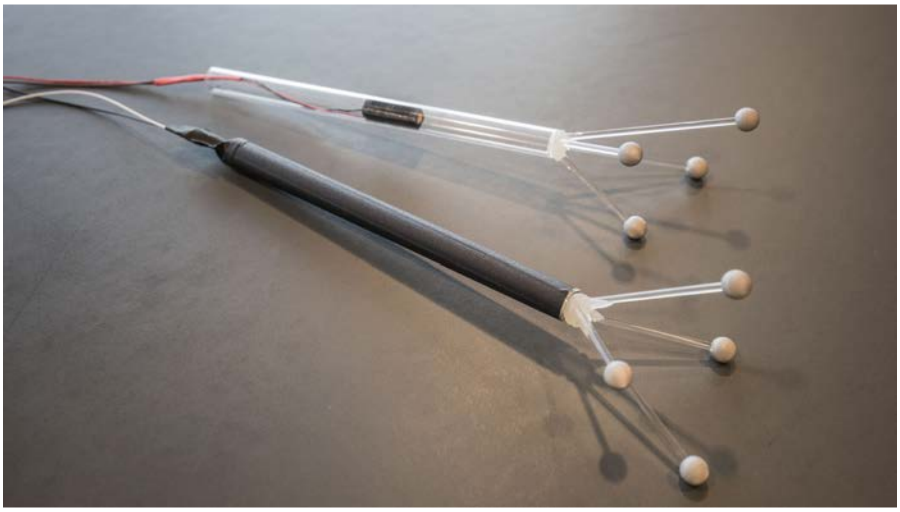

Figure 42: Haptic feedback device used in our study (front, black) and transparent version with position of haptic actuator visible (back, white/transparent). On right side they terminate in IR markers for the OptiTrack system and on the left in a cable that transports the control signal.

a 320 Hz pulseTrain[1]). When the interval used for the pulseTrain is varied based on user motion, an experience of texture can emerge [40, 146, 170]. Based on the physiology of texture perception it may be feasible to generate such an experience in mid-air [13]. To implement such a system, one needs to decide which parameters of user motion to use for controlling the pulse interval. We therefore created a pointing device with a tracking system that allows us to implement various of motion-to-feedback mappings:

10.3.1 *Pointer*

We built a custom pointing device, inspired by the controllers used for VR systems such as the HTC Vive, Occulus Rift or Hololens[2]. We use the Haptuator Mark II[3] by Tactile Labs to generate the vibrotactile feedback. The haptic actuator was placed on the inside of an acrylic pipe, equidistant from both ends. The acrylic pipe had a flexible litz cable on one end, connected to the output of an audio-mixer, and had four markers attached on the other end which were used for tracking its position and orientation (Figure 42).

1 https://developer.oculus.com/documentation/pcsdk/1.9/concepts/dg-input-touch-haptic/
2 See also http://engadget.com/2017/08/25/microsoft-hololens-wand-patent/
3 http://tactilelabs.com/products/haptics/haptuator-mark-ii-v2/

10.3.2 *Tracking*

We measured the position of the pointing device using an Optitrack motion capture system. We use 8 cameras which captured the position and orientation of the device at 125fps. After calibration, the average error in positioning is <1.6mm.

10.3.3 *Signal Generation*

Using Max/MSP, we generate our signal as a pulse train similarly to previous approaches to haptic texture generation [170, 171]. Each pulse has a duration of 1.45 ms (64 samples at 44.1 kHz sampling frequency). The frequency at which they occur is determined by the motion performed by participants, where a fast movement of the pointing device generates a higher number of pulses and holding it still produces no pulses.

10.3.4 *Signal Path*

See also Figure 51, in Chapter 11 for a picture of the full experimental setup.

Position information is calculated by the motion tracking software[4] and passed on to a custom C# application that generates movement information, according to the mapping condition. The C# application sends the movement data to a MAX/MSP patch using OSC. The MAX/MSP patch generates the pulse trains as an audio signal [170, 171]. We used the UR44 audio-interface by Steinberg for signal output to an audio mixer. The audio mixer was used to amplify the signal to the necessary levels for driving the Haptuator, as well as for easily switching between textures. The output of the mixer was connected to the haptic actuator embedded in the pointing device. We estimate the system latency to be < 25 ms[5].

10.4 MOTION-TO-VIBRATION MAPPINGS

An alternative explenation and visualizations of the mappings can be found in Chapter 11

Previous studies used motion-to-vibration mappings defined by the properties of the experimental devices, for instance, sliding [170], pushing [90] or bending [171]. In contrast, when moving an object in mid-air, there is no implicit mapping. We therefore designed three different mappings to understand their influence on the perception of motion-coupled pulse trains .

PROJECTION Rather than a finger moving over a stone wall in front of us, the Projection condition explores the idea of touching a wall

4 http://optitrack.com/products/motive/
5 Camera Shutter Speed: 3.9 ms, Sampling Rate: 8 ms, Networking: 0.85 ms, Motive: 0.7 ms, Max/MSP: <1.5 ms, UR44: 5.12 ms, C# 0.5 ms (Values based on datasheets where available, otherwise measured or calculated)

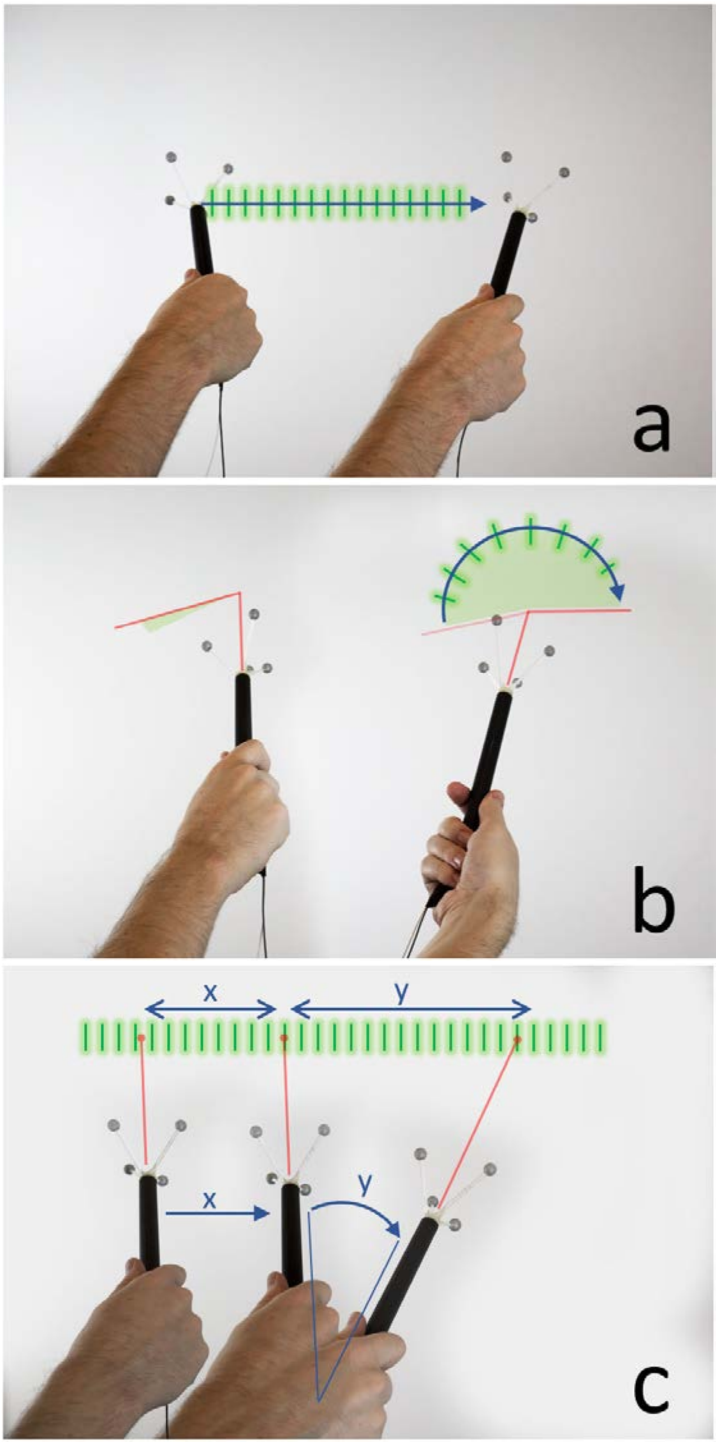

Figure 43: (a) Translation condition (vibration is mapped to displacement of the object). (b) Rotation condition (vibration is mapped to change in orientation of the object). (c) Projection condition (vibration is mapped to change in position of a virtual point moving over a surface). Textures are green, motion is blue and pointer extensions for clarification are red.

that is far away. It is inspired by the light point of a laser-pointer. We generate vibration based on the movement of an imaginary point over a faraway virtual wall. Displacing the pointer (Figure 43 c, marked as ‘x’) generates a steady stream of impulses. Rotating the pointer (marked as ‘y’) causes the imaginary ‘light point’ to move increasingly faster, resulting in an accelerating succession of pulses (Figure 43 c and Video Figure 1c at 00:59). Projection can be broken down into a translation and a rotation component, which we explore individually:

TRANSLATION This mapping is the one most similar to the movement we make when exploring a physical texture and closest to previous work in this area [40, 146, 170]. We measure the position of the pointer in 3D space and map the distance between the objects current position and its previous position to pulse frequency (Figure 43 a and Video Figure 1a at 00:35).

ROTATION We measure the orientation of the cursor (pitch, yaw and roll, as could be sampled from the inertial sensors of smartphones) and map the change in angle to pulse frequency (Figure 43 b and Video Figure 43 b at 00:47). This mapping can be implemented using the IMUs of many existing devices.

10.5 STUDY DESIGN

Our goal was to better understand how mappings influence the perception of motion-coupled, non-grounded vibrotactile feedback. We chose an interview-based approach to ensure that we cover the breadth of experiences people had when interacting with this type of feedback.

10.6 CONDITIONS

We used four study conditions, the first three corresponding to the three mappings explained above: Translation, Rotation, and Projection. In the fourth condition, Vibration, the pointer is actuated by a constant pulse train. For each condition we also presented and discussed the absence of the vibration with the participants (Video Figure 1 at 00:35). We passed this signal through a bandpass filter with a center frequency of 125 Hz and a Q of 250. The low-pass filtering ensured that the vibrations were not audible, and the high-pass filtering made the signal feel crisper. The Vibration condition pulsed at 40 Hz. While we could not control for frequency, due to the different mappings, all mappings were designed to feel as similar as possible[6]. There was no visual or acoustic interface. Participants sat in an ergonomic, armrest-free stool, facing a

6 Note that this is not true for the video figure, where we aimed at making the differences between conditions as clear as possible.

white wall (as seen in Video Figures.) The only information provided to the participants was what they felt in their hands.

10.6.1 *Interview Method*

Our interview method was inspired by Petitmengin [138, 139], with the intent of eliciting descriptions of introspective, subjective experiences. This approach has also been used in a previous study of haptic perception [128] and a study of the 'rubber hand' illusion [183].

We told participants that we research the perception of vibrotactile feedback and that they would be presented with a pointing device that would be vibrated with four patterns. We asked participants to "explore what the pointer feels like by moving it". Participants did not receive explicit instructions on what movements to make. They were asked to maintain the same grip on the pointer for all conditions and, as best as they could, ignore any assumptions they might have about the technological setup and instead focus on their subjective experience.

The interviews were structured by the four conditions which were introduced in rising order of complexity, starting with the Vibration condition (no mapping) and finishing with the Projection condition (non-linear mapping). Translation and Rotation were alternated in order. This allowed participants to slowly build up their own vocabulary, which we then also used when asking questions. Our goal was to explore the breadth and depth of subjective descriptions. Human vocabulary for discussing haptic experiences is limited and initial testing suggested that allowing participants to explore the complexity at their own pace helped them find nuanced ways of expressing themselves.

While conducting the interview we introduced as little information as possible in our questions, using the participants own vocabulary wherever possible. We started the discussion by asking what a vibration pattern felt like and then would follow up by asking participants to expand on their descriptions. If participants made an observation, and inquired if their observation is correct, we always agreed with their observation while asking them to reflect on it further. The following excerpt illustrates a typical exchange:

Exp: I'm going to bring in this third pattern.

P7: (Pause) - Oh, now I can feel that it's responding to how I'm moving it. The vibrations.

Exp: What does that feel like?

P7: It feels quite exciting, actually. I don't think I've ever felt this before ... it feels as if there is something invisible, [...] some kind of force-field that I cannot see influencing it, which kind of confuses my brain a little bit.

Participants were asked to compare all mappings, and to compare them to the absence of haptic feedback. Participants were also asked to compare their behavior and the precision with which they moved the pointer, with and without haptic feedback. Otherwise the topic and pace of the interview was dictated by the participants—we would merely ask for explanations, clarifications or additional elaboration. If the conversation dried up, we switched to the next mapping condition or to one of the predefined questions. We explicitly told participants that they could ask us to switch back to previous mapping if they needed the experience to better make comparisons. Interviews were audio-recorded.

10.6.2 *Participants*

We recruited 12 participants, of which 5 were female, through word of mouth and a university e-mail list. Ages ranged from 21 to 65 years (M = 30, SD = 10.7). All but one participant had completed a university degree. Participants received presents as thanks for their participation. The value of the presents corresponded to a typical hourly wage at the location of the experiment. The interviews lasted between 22 (P10) and 72 (P12) minutes (M = 44.8, SD = 13.7). We initially conducted eight interviews and did a preliminary analysis. We then added four more participants. As no new topics emerged, we decided that the number of participants was sufficient.

10.6.3 *Analysis*

We transcribed the interviews and manually searched them for relevant sections. We discarded statements by the experimenter (except when required for context), and off-topic discussions. The rest of the documents were split into discrete statements and labelled with the participant ID and condition. The interview transcripts had between 2,000 and 8,500 words, totaling about 50,000 words. For reference, this paper is about 9,500 words.

We clustered the statements ‘in vivo’—categories emerged during this process. Data was viewed by all three authors, decisions were made by consensus. We conducted two rounds of clustering. Initially, three major thematic groups emerged. Statements within these groups were then further analyzed and clustered into sub-groups. If a statement fit into more than one sub-group, we created a copy of it, keeping track of duplicates. The clusters emerged by examining how participants responded to the question ‘What does this feel like?’. If they answered by describing the pointer, we placed the response in the object description category. If they answered by describing their actions or what they themselves felt like, we placed the response in the self-observations category. Responses that made higher level observations such as commenting on the process they went through when experiencing the haptic feedback where grouped as meta descriptions. Note that not all topics

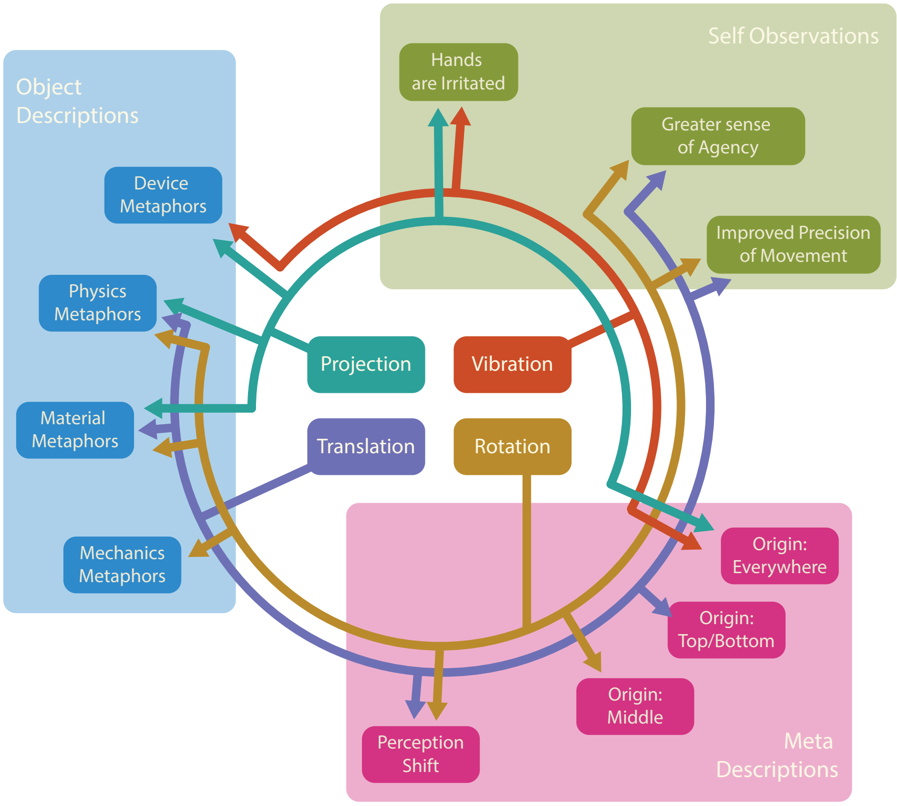

Figure 44: Overview of results. Interview statements were clustered by Object Descriptions (top, blue), Self Observations (bottom right, purple) and Meta Descriptions (bottom left, pink). The arrows point from condition towards main topics and findings. In blue we see metaphors used, in pink we see which conditions led to perception shifts and where users felt the vibration came from. In purple we see effects the experience had on the users. Rotation and Translation share many topics, but here is no overlap in main topics between them and Vibration. Projection shares properties with all conditions.

were covered by all participants. Participants demonstrated very diverse ways of discussing the experiences, evidently drawing from their individual backgrounds.

10.7 RESULTS

As can be seen in Figure 44, there were three main clusters. The four mapping conditions lead participants to discuss these in different ways. From Self Description we learn which mappings were bothersome and which ones might help perform a task. Object Descriptions show us how the experiences of the mappings differ qualitatively from each other and Meta Descriptions teach us about how participants explained what they felt and the order in which experiences occurred. The following is a summary of the 12 interviews.

10.7.1 *Object Descriptions*

When describing what the pointer felt like, participants commonly used metaphors or comparisons to familiar experiences. If these descriptions referred to a material (e.g.: "...there is a ball in the very old mouse for computers [. . .] they had this rubbery surface", they were grouped as 'material metaphors'. Statements describing forces or using other physics concepts such as "the resistance increases if I move it quicker" were labelled 'physics metaphors'. Descriptions of interactions between objects or literal mechanical concepts such as "like when you ride a bicycle and its going too fast for the gear to keep up" where labelled 'mechanical metaphors' while descriptions that referred to electronic devices or electronics (e.g.: "It reminds me of a Geiger counter") were labelled as device metaphor.

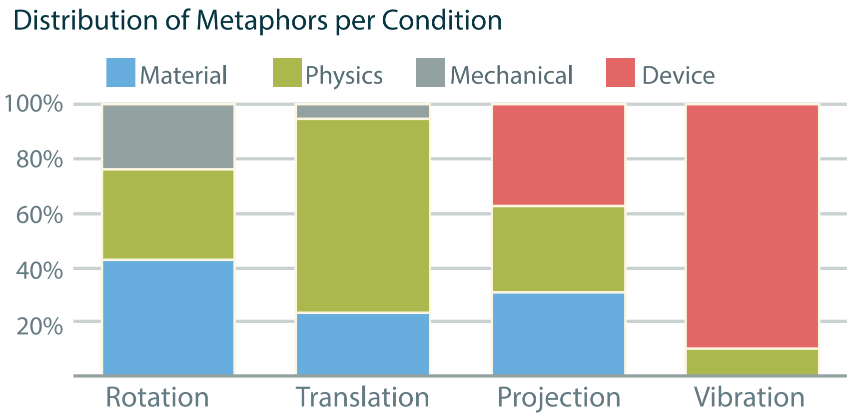

Figure 45: Frequency with which types of metaphors were used in discussing each mapping condition.

Plotting the frequency with which these categories occurred (Figure 45) provides an overview of how participants discussed the mappings. During the Translation condition participant's descriptions mainly used physics metaphors. The Rotation condition elicited the most descriptions of materials, but less discussion of physics than the Translation condition. Instead a large portion of these discussions focused on interlocking gears and other mechanical constructions. P8 explained this difference by stating that they both feel very 'familiar', but that the Rotation condition felt more like something they would expect from a man-made device, while Translation felt more like something they could experience in nature. The Vibration condition was most commonly described by referring to electrical devices (electrical toothbrush, smartphone). The Projection condition combined properties of Translation and Vibration. While participants could clearly distinguish between the Translation and Projection conditions, they grouped them together, often referring to them simultaneously in their descriptions. Participants felt these conditions were more engaging (10 of 12), fun (P1, P5, P11, P12), interesting to move (P6, P7, P8, P9, P10, P11) or pleasurable (P7). For a less fragmented description of each mapping, the rest of this section is organized by condition, not by metaphor.

ROTATION Participants enjoyed moving the pointer with the Rotation mapping. P11 said that "in the beginning [without the feedback] it felt a little boring. [...]. Now it's way more fun". Participants stated they experienced resistance when rotating the pointer (8 of 12). In addition to resistance, participants used terms such as 'higher traction' (P1), and being 'hard to move' (P3) to describe this. The motion also was described as 'grinding' (P11), 'rolling over a rough surface' (P12), and having additional inertia (P11). When asked why this was the case, users explained it by describing the pointer as 'rusty' (P1), 'old' (P1) and 'sticky' (P3). The pointer was also described as feeling heavier (P5, P6, P9). Rolling the pointer in one's fingers was compared to rolling a hexagonal pencil over a table (P12, P5, P6). P6 made the observation that comparing the Rotation mapping to the absence of haptic feedback was like "the difference between a high-quality pen and, um, like a plastic pen" When asked to describe what the haptic feedback made the pointer itself feel like, all users would describe the material composition of the pointer in some way. The most detailed description was provided by P2 who explained that it felt both heavier and softer "[as] if this was made of Styrofoam with like an iron rod or something inside it, to make it a little heavier". P3 and P12 also experienced a certain level of compliance, associating it with rubber. The rotation mapping seemed to make the perception of the pointer more complex. This became apparent from the many multi-material and mechanical metaphors used. P1 and P2 both associated the experience with gear-systems in a bicycle and P3 described it as similar to feeling the rubber ball on the inside of an analog computer mouse. This additional complexity however felt familiar, for example P1 stated "I feel like I've felt this before, but I can't remember where".

When we removed the haptic feedback, users felt that the pointer became 'lighter', as if (P1) 'a gear system loses traction'. P2 also reported that the pointer felt 'colder'. While P11 described it as an additional property of the object that is lost "I can't help but think about it as 'something else' when it's vibrating. It could be 'whatever' in my head. Like a key turning. And then, when the vibrations go away, my imagination fades as well" (P11). When asked if the pointer felt the same in the absence of haptic feedback as it did at the beginning, before they had experienced a mapped vibration, P11 declined, explaining that "it feels like I lost something more than I gained something before. If that makes sense ... Like in the beginning [without vibration] it was alright just to turn it, but now I ... it feels like its missing something when it's not there".

TRANSLATION As with the Rotation condition, participants seemed to enjoy this mapping a lot. P5 and P6 explained that the pointer felt magical or powerful, as if it were a wand from Harry Potter, and P4 immediately exclaimed "This is a lightsaber!". Asked what the Translation mapping felt like, P4 described "So, to me, it's not a vibration… it becomes something else. … it just becomes a resistance, you know, to my movement." Almost all participants (10 of 12) confirmed this experience of resistance. Additionally, P2, P5 and P6 remarked that it felt heavier. P5 specified that "the movements are causing that I can feel that it's heavier". We were somewhat surprised by P3 who felt that the haptic feedback made the pointer feel warmer and by P2 who described the pointer as colder without the haptic feedback.

P4, P5, P7 and P8 described that moving the pointer was as if one was 'moving through a medium'. Examples included 'stirring a pot of dense soup' (P4), 'moving a stick through honey' (P3), and 'swinging a badminton racket' (P2). P12 and P11 experienced such motion as slower, while P3 felt the movement was faster than expected and somehow stickier. P11, P2 and P5 experienced a counter-force when moving it. Compared to the Vibration condition, the experience during the Translation condition was 'cohesive' (P2) or 'more natural' (P12, P2, P3, P7, P8), "because it corresponds to my everyday experiences. It corresponds to feeling something when I brush over it, when I am also moving" (P12).

Removing the haptic feedback presented many of the participants with a feeling of loss, P2 described the pointer, once the vibrations had been removed, as "colder, deader, and lighter". P12 said that "It's like [the experience] is over, because you've put down the object; because it's dependent on your movements" P12 contrasted that to the Vibration condition in which the experiences "gradually fades out".

PROJECTION Many users initially did not experience this as 'natural' in the same way as the two previous mappings. P12 immediately explained that it felt 'artificial' and that what they experienced was 'too complex to correspond with anything natural'. P2 described it as being 'electrical rather than organic'. Similarly to the Vibration condition, participants resorted to metaphors involving machinery. However, this time the vibration that they felt was not a side-effect of the mechanical motion as it was for the Rotation condition, but as output from a digital device—hence we did not consider these in the Mechanical sub-theme. The vibration was

describes as feeling like a 'metal detector' (P4) or 'Geiger counter' (P12). Exploring the Projection condition was often described as a 'spatial' experience (P1, P2, P3, P12). P12 described that it's "complex, because it doesn't correspond to a surface but to the Space around me. When I move it further away or closer to me, with this speed, it kind of corresponds to a spatial experience". As participants felt that the vibrations were caused both by their actions and by their surrounding space, the Projection condition was associated with 'less control' (P1, P2). Participants again were very conscious of the vibrations, P1 stated that it effected their fingers.

VIBRATION Unsurprisingly the Vibration condition was experienced as, well, vibrating. Asked what it felt like, P1, P4, P8, P7 and P12 provided examples of devices that either vibrate themselves ("It's the same as when I use my electric toothbrush. It sort of... tickles a bit" - P1) or objects indirectly vibrated by remote machinery ("As if I would be sitting in the subway and the seat vibrates." - P12). Others (P8, P10, P2) suggested that the vibration felt electric, like "grabbing an electrical fence that is not very high voltage." (P2). Finally, P5 felt said "well my first thought was, that it was a bit stressful, I guess. Like that I should um, like if you get a notice or an alert or something" While these results appear rather obvious they provide a useful contrast to the other conditions

10.7.2 *Self Observations*

When participants answered the question 'What does it feel like?' with a description focused around themselves, we placed the description in the self observations category. Within this category two groups emerged. One, holding the pointer consisted of descriptions of how the haptic feedback influenced their hands. The other, moving the pointer was participant's descriptions of their behavior when interacting with our system.

HOLDING THE POINTER While vibrations can have a positive, relaxing effect (e.g.: [20]), many participants, however, (7 of 12) commented negatively on the experience of holding a vibrating object. Participants remarked that the vibration interacted with their fingers in an unpleasant way (P1, P2) i.e.: "If I did this for a while it would feel like my hand was a sleep" (P2) and that this unpleasantness continued even after the vibration was removed "in my hands it still tickles a bit" (P3).

Of all people who commented negatively on the vibrations most (5 of 7) pointed out that the unpleasant feeling went away for the mapped vibrations of the Translation and Projection condition. For example, when describing the Translation condition, p3 stated that "It did not spread out that much. It felt more like that the resistance was in [the tip of the pointer]. It's not that I didn't feel it. It was just ... my hand is fine now. I can't feel it now. But after [the vibration condition], I could still feel it after it was out, in my fingers. [The vibration condition] kind of left traces."

MOVING THE POINTER Participants stated that they felt the Translation and Rotation condition provided them with a more exact understanding of their movements. While they did not believe that they could move with more accuracy, they felt that they could reproduce a movement with more precision (9 of 12).

Asked to draw an infinity sign and given the option to use any of the conditions, participants typically (9 of 12) chose the Translation condition. P2 explained that "[If I had to draw] one or two perfect infinity signs, then it probably would not help me so much, but if my life depended on drawing a thousand, then it probably would". This heightened sensitivity to their movements also influenced their behavior. Typically, participants moved the pointer slower with haptic feedback present and faster without it. When asked why this was the case, P12 explained that "It's because I was paying attention to the impulses before, as they were reacting to my movements. Now, without the impulses only my motion is left without the additional perception I had before".

10.7.3 *Meta Descriptions*

A third way of answering the question 'What does it feel like?' were responses that took a broader view and attempted to contextualize the experience, either in time or in relation to others. The two most important themes from this grouping were descriptions of the origin of the vibration and descriptions of a perception-shift, from experiencing a pulse train to a richer material experience.

ORIGIN OF THE VIBRATIONS In the Vibration condition, participants did not have any specific impression regarding the origin of the vibrations. They just seemed to come from the pointer. This was different for other mapping conditions. P1 and P3 stated that they felt that for the Translation mapping the vibration came from the top of the pointer. P5 felt their perception switch between top and bottom while P2 felt it came from both ends at the same time. This was related to the experience of P8 who felt that the vibration was caused by the environment the pointer was in and P5 and P7 who stated that the vibration felt like it was caused by a medium the pointer passed through. P2 specified further that while the vibration comes from the top or bottom, the pointer is vibrated from its core. P12 also felt that the vibration came from the core of the device.

The origin of the vibration was less clear for the Rotation condition. Here participants did not feel that it came from the ends of the pointer. P5 for example felt that it was 'everywhere'. P12 and P2 felt that the vibration originated from the surface of the pointer. This is contrasted by P11 and P10 who felt that the vibration came from within. P10 stated that the vibrations felt "as if there were an object inside the pointer".

In the Projection condition P12 and P11 felt that the vibration was caused through interaction with the room. P12 felt that the origin of the vibration could be either close or far, depending on the mental images used to think about it.

PERCEPTION SHIFTS One of the most interesting questions was how something that was considered to be an irritating vibrating object could transform into something with new physical properties that participants enjoyed moving. Throughout the interviews we were able to identify clear steps in this process:
(1) Initially participants would focus on the impulses themselves, often unsure what they are experiencing. For example, asked to describe their experience, P8 explained that "I don't know what the right word for it is, but it, they're sort of discreet pulses" while P11 was displeased that it was "just vibrating every time I put it anywhere". P7 also disliked her initial experience, stating that "it kind of feels like it's in pain".
(2) Eventually participants learned to understand the mappings and focused their descriptions on them, sharing observations such as "When I move the stick here, its vibrating and when I stop it stops as well" (P3) or "Ahhh so it vibrates faster as I move it faster and it vibrates slower when I move it slower" (P7).
(3) Understanding the coupling between motion and vibration was not sufficient for creating material experiences. Instead it appears as though at some point, when interacting with the mapped vibration, the perception somehow shifts. For some this happened very fast, others had to move back and forth between couplings for this shift to happen. Sometimes we could tell, based on exclamations, that the shift had occurred. Participants spontaneously exclaimed the following, after having experimented with the mappings: "Oh, that's neat. It's just so different, it, it, I mean, that reaction, it's totally different" (P7) or "Woah. It feels like coloring with vibrations" (P3) and, possibly inspired by the magic-wand like shape of our pointer, "I have magic power, I think. (laughs) I don't know. Yeah, magic power" (P5).

Once the perception-shift occurred, participants no longer described the feedback in terms of impulses or couplings, but in terms of interactive experiences. This perception-shift lead to an experience which was both qualitatively different than before as well as novel. For example, P11 described the process of adding a mapped vibration to the pointer that "It became more" and that "the vibrations [...] make it feel like something different, like something it's not".

We observed an interesting tension because of the way the perception-shift lead to an experience that participants were both familiar with from the physical world, while at the same time being very foreign. This is captured in a description by P8: "I have this, uh, 3D printer that has a little knob that you use to [control a] simple interface, and it kind of clicks like that. So, that was actually what came to mind. But it, it's kind of, it's not something you're used to experiencing". The foreignness of the experience was also described by P7 who, when their perception of the Translation condition shifted, explained "it feels quite exciting actually like this. I don't think I've ever felt this before [...] it feels as if there is something, like there is something invisible, like obviously there isn't but ... but there is some kind of force field that I cannot see, influencing it. Which kind of confuses my brain a little bit". The dissonance between what participants experienced and what they thought they should experience is also highlighted by P7 who later expressed her worry: "I hope you don't think I'm crazy".

The perception change was difficult to achieve during the Projection condition due to the complexity of the mapping. Only two participants (P8, P12) were not able to create a mental model that fully explained the mapping. Most others felt that there were external sources influencing the feedback, which introduced a source of uncertainty, preventing the perception-shift from fully establishing itself.

10.7.4 *Breakdown Conditions, Limitations*

The Translation and Rotation mappings lead to very strong material experiences. However, they required users to move both 'correctly' and very steadily. For example, if a participant changed the position of the pointer in the Translation condition they would almost always also change its orientation. Similarly, it is hard to change the orientation of an object in free space without also slightly changing its position. When the movement and the mapping did not match, participants felt vibration that seemed to react to their movements without synchronizing to what they were doing. In these situations, users commonly considered the pointer as if it was a small creature, with intentions and agency of its own. For example, P3 described this as "an animal sleeping and just moving a bit around. It's not much, just very little vibrations. So, for example, there is nothing now, but if I move my hand to roll it around, it moved."

When we touch a physical surface, we can move our hand very steadily, because it is supported by the material we are touching. In midair, such steady motion is difficult. Sometimes user's hands move slightly without the user intending it. This caused impulses which did not match the expected behavior of the pointer. Both these issues of mismatch between motion and mapping, as well as user precision pose a design challenge.

10.8 DISCUSSION

There is a perceptual link between movement and vibration. Together, they enable us to experience textures. Various prototypes build on this observation to create virtual textures. As these virtual textures can be created without the normal force of a supporting surface, we can conceivably render textures in mid-air. Mid-air textures do not have a clear motion-to-vibration mapping as in other haptic rendering systems. It is not even clear, what the properties of a 'good' motion-to-vibration mapping should be. Therefore, we set out to explore various motion-to-vibration mappings. Two of the mappings we tested, Rotation and Translation, did indeed have properties that allow us to consider them as 'successful' mappings. Participants experienced additional friction, force and weight when interacting with them. While we do not suggest that these sensations are identical to the texture experience when touching a surface, based on the interviews we find them sufficiently similar to use the same word for both. From now on we will use the word 'texture' in a loose definition that includes mid-air textures.

10.8.1 *Benefits of Using Textures*

We found that textures had a series of benefits over using regular vibration as haptic feedback:

LESS IRRITATION Using textures for conveying vibrotactile information does not cause the irritation that is often associated with vibrotactile feedback. Textures are interesting to move through, while regular vibration is considered irritating in a similar way that an electric toothbrush might be.

ADDED 'QUALITY' Adding textures to an object changes its perceived quality, when moving objects through a texture they were perceived as heavier which was experienced as being 'higher quality'.

PROVIDING USERS WITH AGENCY Using textures provides users with a sense of control. If information is represented as regular vibration, the user has limited control over when or how it is perceived. Using mid-air textures provides the user with a way of anticipating what they should feel if they perform a given action. Information can be encoded to change the 'feel' of that action. The user can then explore this information at a time and pace of their choosing. This provides user with a greater sense of agency over the interactive system.

IMPROVING THE EXPERIENCE OF CONTROL In addition to providing users with a sense of agency, users also feel as though they can be more precise in their movements. The texture acts as an additional feedback channel that helps users observe their own actions. Users felt that textures help them repeat the same gesture multiple times with higher precision. This could be used to provide people with more confidence in using gestural interfaces or improve movement and gesture learning.

This was something users also described to us when using ReFlex *(Chapter 5).*

10.8.2 *Concrete Applications using Textures*

Here we present some simple interactions that leverage these textures to different degrees (see Figures 46 – 50 and Video Figure 2). The intent of this section is to demonstrate how existing gestural interaction systems can leverage our results.

HAPTIC TARGETING Haptic systems are often used to convey spatial information. A device might vibrate when pointing toward an interesting location [144] or the pulse frequency or timbre can be modulated based on the distance to a target [51] (See Figure 4a and Video Figure 2a, at 01:25).

We created such a targeting application. We created a texture using the Translation mapping and modulated timbre and amplitude based on the distance to the target. We also implemented a version without the mapping. Anecdotally, users are able to find the target in both versions, but the textures were more pleasant to interact with than vibration.

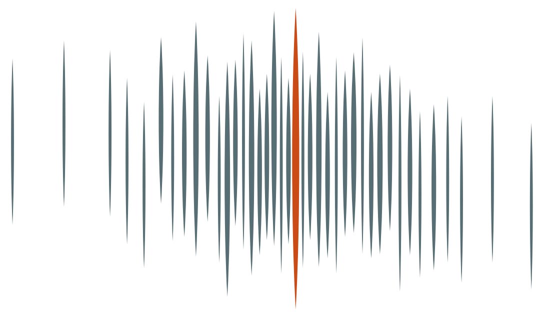

Figure 46: Haptic Targeting: Amplitude, granularity or a specific timbre might be concentrated in an area, indicating the vicinity of a haptic target.

DIRECTIONAL HAPTICS One of the problems with haptic Targeting was described by user P5 who stated, "I feel like I'm looking through a periscope". What they were referring to is true for both textures and vibrations. It is difficult to get the 'big picture', as one can only experience the singular contact point between the pointer and the virtual texture. This makes finding haptic targets a chore that involved carefully scanning through volumetric space. We created an application that provides users with a sense of the direction of the target (see Figure 47 and Video Figure 2b, at 01:37).

By tracking changes in distance, we know if the user is moving towards or away from a target. We can generate a different texture based on the user's movement. The directional haptics appeared to primarily help find the general vicinity of the target. We suggest this approach be combined with haptic targeting for finding the precise location.

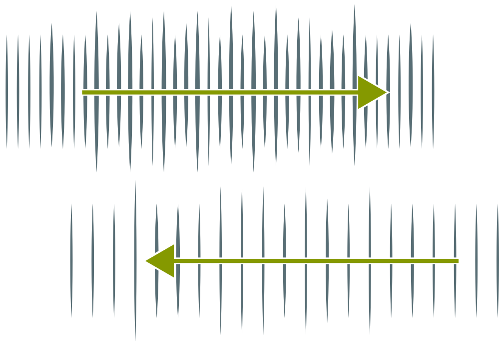

Figure 47: Directional Haptics: Textures can be designed to change based on the direction one moves through them

DIRECTIONAL OBJECTS The Projection condition, which provided spatial information to users, was both the most confusing and the least enjoyable mapping. This suggests to us that providing spatial information may not be the application that best takes advantage of textures.

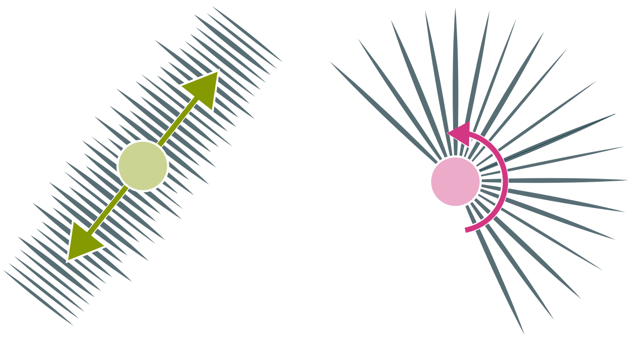

Figure 48: Directional Objects: The green object communicates haptically that it is mapped to linear control, the purple one to rotary control.

Rather than using mid-air textures for indicating locations in space, we can use them to provide objects with additional affordances. By constraining the mapping to a single dimension, or a small number of dimensions, we can give objects 'directionality'. If we map a texture to the 'roll' dimension, a movement in this dimension sticks out relative to other movements. We use this phenomenon to create a series of controllers with prescribed mappings. (See Figure 48 and Video Figure 2c, at 01:49).

We created a slider that could be controlled by movement in a single arbitrary dimension (either x, y, z, pitch, yaw or roll). Upon picking up the controller, one can immediately identify the required motion for adjusting the slider without requiring a visual aid.

DYNAMIC OBJECTS Such augmented directionality need not be static. The 'direction' of the object can be changed when context or tasks switch. In addition to changing the 'direction', the scale (by adjusting granularity, see [170]) can also be modified. This can convey to the user if they should perform a fast or slow gesture (See Figure 49 and Video Figure 2e, at 02:04).

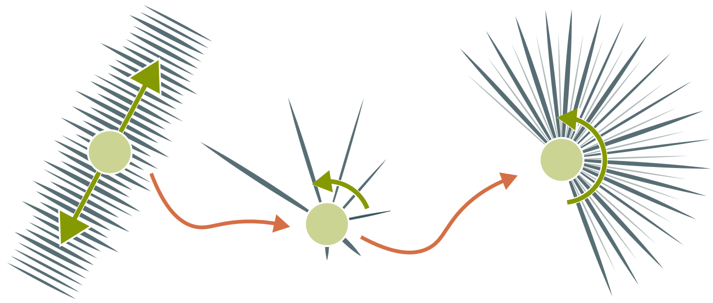

Figure 49: Dynamic Objects: Based on context the green object might switch functionality from slider to rotary encoder. The granularity can indicate the required precision of the application

AUGMENTED PROPRIOCEPTION Users feel that they can move with greater precision when moving the pointer through a texture. The perceived ability to move with greater precision would appear most useful if one's hands were free to manipulate one's surroundings. Moving the vibrotactile actuator from the handheld pointer to a wearable device preserved the experience of precise motion while allowing the user to hold tools or perform gestures unburdened by a tangible pointer (See Figure 4f and Video Figure 50, at 02:17). While worn on the wrist, the feedback can still provide dynamic directionality, suggesting preferred movements to the user.

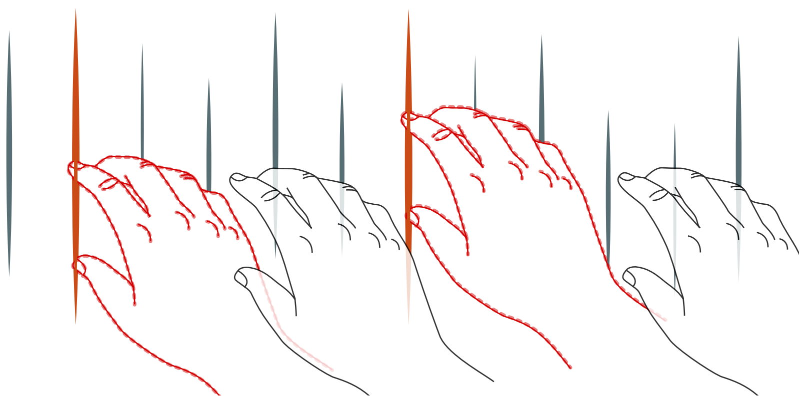

Figure 50: Augmented Proprioception: When the hand moves through space, the user perceived pulses at fixed intervals. Participants felt that this increased the precision of their movements. The pulses could be delivered using a haptic wristband or a system such as *Magnetips* (Chapter 2) [115].

10.8.3 *Key Conceptual Takeaway*

I expand upon these ideas in Chapter 14.

For us, one of the most interesting observations was the clear perception-shift from the participants' hands being vibrated by the pointer to the participants experiencing textures through the pointer.

This switch can be considered a Gestalt phenomenon: Individual pulses are bound together by movement and perceived as a larger interaction gestalt [165, 175], similarly to how the black shapes of Kanisza's triangle (see Figure 63) lead us to perceive a white triangle [57]. Like bi-stable images or foreground-background illusions the emerging texture experience is also multi-stable. Users could change what they perceived by 'imagining pictures' or 'changing their intention'.

Another way of thinking about this perceptual shift is one of attention and agency. Participants initially focused on the object and haptic information provided to them through the object. When experiencing a texture, the focus of attention moved beyond the pointing device, the attention was directed at the interaction. The pointer transforms from an object that is being observed, to a tool through which participants actively explored the haptic experience.

The perspective switch changed the way information provided through vibration is interpreted. Before the perception switch, vibrations are experienced as symbols that provide information, for instance, when operating a telerobot, a user might receive a vibration, symbolizing that the robot is being touched [164][7]. In our case, once the perception switch took place, vibrations were no longer considered as symbolic carrier of information. For example, when participants stated that moving the pointer felt heavier, they did not mean that the vibration represented 'heavy', they meant that they experienced increased weight.

10.9 CONCLUSION

I use the term "material experience" as a catch-all for the various experiences which can arise based on mapping and feedback parameters.

So, can we experience haptic textures in mid-air? We found that, based on mapping, experiences very similar to texture can be created. If an object just vibrates without reacting to movement, it is experienced as a device, such as a toothbrush or vibrating smartphone. If an object vibrates based on where it is pointed, it feels more useful, but still like a device – maybe a Geiger counter or metal detector. In our Rotation and Translation condition, however, the way the vibrations were experienced transformed, leading to a material experience related to texture. These textures are more pleasing than 'traditional' vibration and make moving a device more interesting – as if it had higher material quality. Systems using mid-air textures can provide users with a stronger experience of agency and a better sense of control when interacting with them.

7 See also Ihde's account of hermeneutic mediation and embodied mediation, as summarized by Verbeek [186].

10.10 ACKNOWLEDGEMENTS

We wish to thank JP Carrascal for his assistance with the analysis. This work was supported by the European Research Council, grant no. 648785.

11 ADDITIONAL INTERVIEW EXCERPTS

The page count of the published CHI paper – *From Pulse Trains to "Coloring with Vibrations": Motion Mappings for Mid-Air Haptic Textures* [166] (Chapter 10) – required us to severely limit the interview excerpts we shared. However, many of our claims are better understood with more context. I will use this section to share extended quotes and dialogues from the interviews. Unlike the paper, were we provide a comprehensive overview, the selected excerpts here are the ones I find most interesting, curious or delightful – from a perspective of some distance since the experiment was run and the paper was written. The excerpts are edited to improve readability, while care was taken to preserve the voice of the participants.

For quick reference a list of the experimental conditions is provided in Table 9.

Vibration	Translation	Rotation	Projection
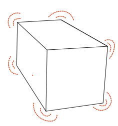	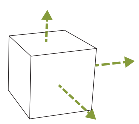	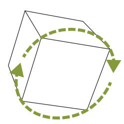	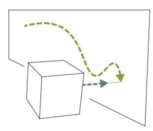
The pointer vibrates at a fixed frequency	The pointer vibrates at a frequency relative to how fast it is moved	The pointer vibrates at a frequency relative to how fast it is rotated	An imaginary ray is emitted from the pointer. The pointer is vibrated at a frequency relative to how fast the intersection of the imaginary ray and an imaginary wall is moved.

Table 9: The four experimental conditions explored in *Pulse Trains* (Chapter 10).

11.1 PERCEPTION SWITCH

11.1.1 *Rotation Condition*

It is possible that this is because the Optitrack system appeared more sensitive to rotations than to translations of objects, making the translation condition feel slightly sluggish compared to rotation.

While we assumed that the translation condition would be the most obvious to participants, it turned out that often the rotation condition lead to the fastest switch. Below are two instances of participants first experiencing the perception switch with the rotation condition.

Participant 6 verbalizes their thoughts after I switch from the vibration to the rotation condition:

P6: Okay, so, so first of all, the pulsation has stopped. I'm now trying to move it in order to get anything. I'm still not getting anything.

So, yeah, there's a slight thing here. Okay. So sometimes I get like two pulses but I'm not sure how to trigger this.

So let's try systematically. X, not so much. Y, nothing, oh ... okay. So there is something but I'm not sure how to get it ... oh, yeah, okay so, so its when I change the orientation of the [pointer], but its not consistent. So I'm not sure. So now the pulses are quite different. It's not constant, but it's not systematic as well so it doesn't relate to how much I tilt or move.

. . . so it's not correlated with the movement. It's more correlated with a change of the orientation. If I rotate it in either direction, okay, now it's more consistent.

Oh. Okay uh . . . yeah, I don't know why it was doing this. It's now much more like a knob, uh with haptic feedback. Maybe like the one in cars when you try to adjust the AC or the volume or something.

Participant 1 describes the rotation condition relative to her previous experience with the vibration condition:

P1: Ok, yeah, thats very different. Thats eh . . . ok yeah, now I can really feel that its my movement that makes the vibrations or whatever it is . . . so I feel much more in control. I can't really relate it to any . . . it doesn't make my think of any object. . . . but I feel like I'm. Like I . . . Like I've felt this before. But I can't remember where. Hm . . . maybe its . . . something like, you know when you bike your bicycle and you finally - if its too fast for the gear to keep up - then when you finally hit it (or how do you say it?) like when you finally reach the point where it starts moving maybe its something like that.

I believe the word P1 is searching for is traction.

11.1.2 *Translation Condition*

Interestingly the translation condition took longer to understand, often requiring me to nudge the participants in the right direction.

Participant 11 holds the pointer and initially experiences nothing. I try and nudge them towards experiencing the shift in perception:

Exp: Have you tried moving it?

Po: Oh . . . no . . . now its just vibrating everytime I put it anywhere.

Exp: Have you tried moving it slowly?

Po: Ahhh so it vibrates faster as I move it faster and it vibrates slower when I move it slower.

Exp: Try closing your eyes. I'm gonna make it go away and I'm gonna make it come back.

Po: Yeah, similar to before, it feels weird when its not there, but, its nice to have this kind of feedback. In another sense. Like adding a new feeling. Its cool with a new kind of feedback, because usually I would not be able to feel it when I'm moving it around.

Exp: Can you describe what you're feeling?

Po: Maybe it feels like more of a pushing through something? Like the air becomes more of a force that I have to . . . if I move it slow, it moves slow but easy, but then the resistance increases if I move it quicker. There is more pressing against it. It gives a feeling of moving something with more resistance than without it.

Participant 7 initially interprets the haptic experiences very differently than I expected. By moving back and forth between conditions, eventually they experience the perception shift:

P7: Well it doesn't feel nice.

P7: It feels like a little bird that's trying to like escape or something. It feels like it's in pain or like (laughs) its struggling to have a heartbeat or to be free. It kind of feels like its dying.

Exp: Why?

P7: Because the beat is irregular. So its like struggling a little bit. I think I anthropomorphize everything but ... it feels a bit more stressful.

Exp: Can you try moving it from the left to the right?

P7: Like this?

Exp: Yeah. Does it still feel like a dying bird?

P7: Yeah. It's because of the beat of it though.

(I ask participant 7 to compare the translation condition to the vibration condition.)

P7: The [vibration] was a lot more pleasurable and relaxing and the [translation condition] was more kind of like uh, music where the beat is like irregular. [The vibration] I find really um, relaxing so I think that the rhythm of the vibration had a big impact on the experience.

P7: And it was very different but I don't think the moving was different. It's almost as if the [the pointer], is different, its got such a different quality, the [pointer], when the vibration patterns are different that the moving and everything doesn't influence that.

Exp: How are the qualities different?

P7: Because of the, regular vibration patterns and the irregular vibration patterns.

(I let p7 feel the vibration condition again)

Exp: I'm going to bring the [Vibration] back - you called this regular.

P7: Regular, yeah. I can feel that its got a beat, a rhythmic beat to it.

Exp: Hmm.

P7: The vibrations feel more relaxing because I can feel it more consistently in my body.

(I switch back to the translation condition)

P7: Oh, now I can feel that it's responding to how I'm moving it. The vibrations.

Exp: What does that feel like?

P7: Its, it feels quite exciting actually like this. I don't think I've ever felt this before ... it feels as if there is something, like there is something invisible, like obviously there isn't, but there is some kind of force field that I cannot see influencing it, which kind of confuses my brain a little bit. I mean I kind of know how the technology is working but like, it's kind of magic a little bit because I know it feels as if there is some kind of stone wall that its, that its grating up against which only responds when I'm moving it and yet there is nothing there.

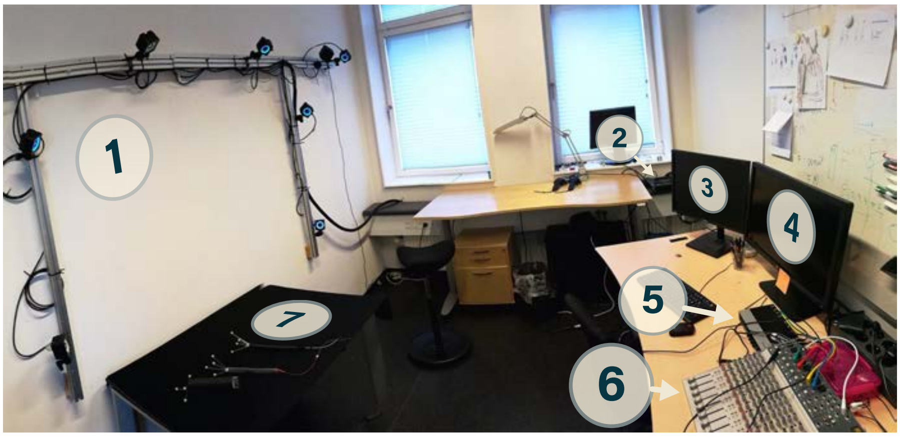

Figure 51: Experimental setup used for the *Pulse Trains* study. (1) Optitrack cameras mounted on wall, connect to (2) ethernet switch. (3) C# application extracts relevant motion information from Motive. (4) Max/MSP patch generates audio signals for all conditions concurrently. They are all played back via the (5) UR44 audio interface and sent to a (6) mixer. The mixer is used to adjust the amplitudes and select signals which are then sent to the (7) haptic pointing devices, which in turn are tracked by the (1) Optitrack cameras.

11.2 TRANSITIONS

A point not emphasized in the paper, which feels important in retrospect was the transition behavior. I could manually fade the different mappings in and out, using an audio mixer. When I did this, the process of fading was experienced differently for the vibration condition and for the conditions where a perception shift occurred. The vibration condition was perceived as gradually fading in or out, while the mappings which created a perception shift were experienced as more dichotomous, either there or not there:

Participant 12 clearly differentiates the fade-out behavior of the vibration condition and the translation condition. P12 traces the difference in experienced behavior to what they attribute the origin of the vibration to be:

Exp: I showed the [vibration condition] to you, and when I took it away, you described it as a "vacuum cleaner gradually being turned to lower power". I'm going to gradually remove [the translation condition]. Can you describe if this is the same?

P12: Its different, because its not gradual, but dependent on its movement.

Exp: But is the way you experience it still the same? Would you still use the same metaphor?

P12: Well, its as if I put something down, and then its gone, because it depends on my movements.

Exp: But is the process the same?

P12: No, its dependent on my movements, depending on what I do. The other one is external, caused by the artificial impulses, which are just linearly reduced.

Participant 3 described the vibration condition as a spaceship, while rotation and translation where complex mechanical structures involving rubber balls. Here P3 describes the difference in fading in and out:

Exp: When I fade the patterns in and out, does the change happen at the same pace or is there somewhere were the change all of a sudden is faster, and you feel like 'oh, now it changed'.

P3: Hm. I think for the spaceship (note: vibration condition) its kind of just continuous, there is no difference. In the [rotation Condition] it's also like first pushing the ball into it and then there is something held to the stick, so I think there most be some kind of difference in how I felt it.

11.3 MATERIAL EXPERIENCES

Experiences of friction or moving the pointer through a medium can be found throughout the other quotes. Here I would like to highlight the concepts of weight and quality, which I personally found especially interesting.

Participant 6 mentions that the pointer feels heavier with feedback that leads to a perception-shift:

P6: Okay. Um ... yeah, that's, yeah, I think weight-wise, yeah, it feels much more heavy than with no feedback. Okay.

Participant 5 mentions that the pointer feels lighter after feedback that leads to a perception-shift is removed:

Exp: So, so without feedback, what does it feel like now?

P5: Now it just feels like a normal object. It's kind of ... I don't know. Em, it feels like it became lighter as well. Like before, it was more heavy. Mm, yeah.

Participant 5 explains the difference in experience as relating to the quality of the pointer:

P6: This is kind of more, more reliable or more sturdy. But the other one is, okay, just with no, with no feedback or nothing. Yeah. So I mean this weight in a, in a good pen or, or in a good something like in a good device that you feel like this is a high quality or a high built ... rather than, okay, this is something that's plastic or something for just doing anything. So yeah.

11.4 CONTROL & AGENCY

Once the perception switch had occurred, participants almost unanimously *liked* the experience of moving the pointer. This was often explained by using the words control or agency.

Participant 3 initially spoke of control *rather than experience, when asked to describe the translation condition. This excerpt is also where* colouring with vibrations *from the paper's title is lifted from.*

Exp: So what does that feel like?

P3: Like ... having control over the movements. ... I kind of want to all the time stop and then start again, just doing it. It feels like colouring with vibrations. Cause you can ... I know there is no difference in the vibration, but I can move the stick in different directions so I feel that I kind of form the vibration.

I asked all participants to draw symbols into the air, and describe how the different feedback conditions influenced the experience.

Participant 8 describes that the textures added constraints, which made them feel more confident in their actions:

(Participant 8 Initially describes how drawing without a texture feels compared to drawing with a texture.)

P8: So, it's, it's almost like the, uh, the rod gets lighter, in a way. Um, and there's no resistance, I think I mentioned before. It's, it was like there was a resistance before. Um, and that's not here anymore. Now, I can just move it wherever I want. Uh, it's much less interesting, but it, it was almost like there was, uh, there was a resistance before, or there was at least something reacting to my movement. So, um, it wasn't just totally free.

Exp: All right. That's interesting. Do you think that in one of these two ways it's easier or hardier to draw an infinity symbol?

P8: Uh, I think the first one was better. Um, when ...

Exp: So, you're saying with haptic feedback, it was better.

P8: Yeah. It was, it was easier, uh, in a way. There's too much freedom without, without it.

Exp: So, maybe you've already answered, but can you describe what makes you say better?

P8: Uh, I, I literally feel like I drew a better symbol, um, although it's hard to judge, but, uh, but it feels like it's, it's, without the haptic feedback, it's, it's easy to, uh, go off course, to go off, uh, uh, to, to make a wrong movement or something else. It might just be an illusion, but that's how I felt.

Participant 7 framed the difference between experiencing a vibration where the perception shift had occurred to regular vibration in terms of agency, rather than control:

P7: Yeah, so like you invite me into an experiment and you give me something to hold and you ask me how does it feel and therefore I go, I put my mind into my body and think okay I can feel it in my lungs a little bit, I can feel it, its making my shoulders relax and then you change the object so that I have agency over it.

Exp: Mm-hmm (affirmative).

P7: So it switches around so that I can uh, and therefore I'm not thinking that much about my body. I'm thinking about the uh, the kind of texture of the experience um, my experience of the object kind of thing.

11.5 SCOPE AND LIMITATIONS

This study was exploratory and I make no claims regarding generalization. Each description is personal and unique. The purpose of this study was to explore the breadth of experiences that participants have with the experimental setup, as well as the themes that emerged in discussion. The original reason for choosing this approach was that a large portion of the related work left me thinking "OK, thats good and all, but what was it actually like?"

The value of such subjective data goes beyond satisfying that itch though. While not generalizeable in of itself, the material experiences discussed by participants provide point of references for future explorations. The descriptions of the perceptual process provide a starting point for an empirical theory of embodied perception in HCI. Using these subjective accounts we can create an initial set of conditions required for the perceptive switche to occur, which can then be refined by further empirical experimentation. A large part of Chapter 14 is based on these interviews.

12 STAND-ALONE PROTOTYPES

As the setup used in *From Pulse Trains to "Coloring with Vibrations": Motion Mappings for Mid-Air Haptic Textures* [166] (Chapter 10) was not mobile, I decided to build a mobile implementation, so that I could easily demonstrate the principles when travelling and attending academic events. This section outlines the implementation. I share it, because it is easier to replicate than the original system.

I built two prototypes. One used rotation as input, the other used pressure. I chose rotation as it is easy to implement using the sensors already integrated in most smartphones. I chose pressure because I was curious how it might work and because I am interested in integrating such systems into the sole of shoes.

12.1 IMPLEMENTATION

Both devices are based on a Teensy 3.2 augmented with a Prop Shield[1]. I used the Teensy Audio Library to generate an AC signal using the onboard DAC. The output of the DAC is connected to the audio amplifier integrated in the prop shield. The output of the amplifier was connected directly to the Haptuator Mark II by Tactile Labs[2]. The devices were powered by generic portable smartphone batteries.

The main difference between the devices is their sensor input. The rotation device uses the Euler Angles of the Prop Shield's IMU. The pressure devices uses a custom-made piezo-resistive pressure sensor, built using 20 kOhm per square piezo-resistive material by Eeonyx[3]. The pressure sensor used two layers of the piezo-resistive material, so that the dynamic range captured the pressure exerted by a foot when stepping on it.

The signals were initially passed through a low-pass filter, to limit unwanted noise. I then selected a threshold for each sensor and once the threshold was crossed generated a pulse-burst at 200 Hz. The threshold corresponds to the granularity parameter, the frequency of the pulse-burst corresponds to the timbre. General principles one might consider when using such a setup in a real world application can be found in Chapter 15.

1 https://www.pjrc.com/store/prop_shield.html
2 http://tactilelabs.com/products/haptics/haptuator-mark-ii-v2/
3 https://www.kobakant.at/DIY/?p=913

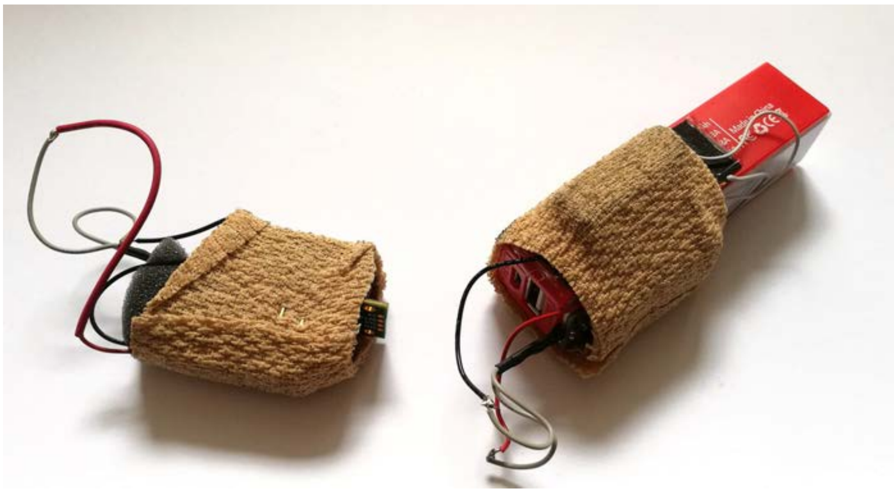

Figure 52: Left: Rotary feedback device. Right: Pressure based feedback device

12.2 REACTIONS

These devices were informally demoed at CHI 2018, at the *In-Touch* workshop[4] [62] and the Dagstuhl Seminar *On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces*[5] [163].

When demoing the devices, I did not tell people what to expect. Most, but not all, were able to make the perceptual switch. A common metaphor for the rotation device was using a rotary combination lock. Pressing the pressure based device made it feel as though it was more compliant than it was. For some this sensation was bistable and could switch to the impression that compression caused internal mechanical motion in the device. Stepping on the pressure based device felt like walking on an old creaky floor.

I had the opportunity to demonstrate these devices to a congenitally blind person and discuss these devices with them. Asked what the devices felt like, they responded, without hesitation "This one pulses when the pressure level changes, and this one provides pulses relative to how it is rotated". When I explained that for many sighted people material experiences emerge, their comment was "Oh, sighted people . . . they are so easy to fool".

4 https://intouchchi.wordpress.com/
5 https://www.dagstuhl.de/en/program/calendar/semhp/?semnr=18212

Part III

IMPLICATIONS

There are things I learned from conducting this research which are not contained in any single paper. In this section I synthesize the results from all papers, highlighting how they relate and how they, together, are more than the sum of their parts. In Chapter 13 I summarize the results of all four papers. In Chapter 14 I highlight how these results and observations made while conducting the experiments are relevant to understanding perception and embodied interaction. Finally in Chapter 15 I describe a technology which I would like to implement. I use this technology to explain how the results outlined in Chapter 13 and the ideas and concepts from Chapters 14 might be applied.

This is a start. It might seem odd to present a start at the end of a thesis, however, I find it natural that research leads not to definitive answers, but rather to new sets of questions. In many ways, this start is where I wanted to be when I began the research leading up to this thesis – however, at the time I was not ready to formulate it. As such this start is not only a hint as to what might come next, but, to me, also a satisfactory conclusion of a period of research.

13

RESULTS OVERVIEW

In this section I summarize the direct results of the studies presented in this thesis, as they are relevant for designing material experiences. Here I only enumerate the pragmatic empirical results, a discussion of their higher level implications can be found in Chapter 14 and an example system describing how the results might be applied can be found in Chapter 15.

13.1 OUTPUT PARAMETERS

When I speak of *feedback parameters* I am referring to the properties of the vibrotactile signal that a user is subject to. Generally speaking these are the strength of the signal (i.e.: *amplitude*), parameters shaping the signal in the frequency domain (e.g.:*timbre, frequency*), and parameters shaping the temporal properties of the signal (e.g.: *granularity, regularity, duration*). Here I summarise various ways in which changing these feedback parameters influences the resulting experience.

13.1.1 *On-Body and Implanted Magnetic Actuation*

I explored direct stimulation of the skin in the case of the magnet attached to a fingernail in Chapter 2 and, using an implanted magnet, I explore in-vivo stimulation of the skin in Chapter 4.

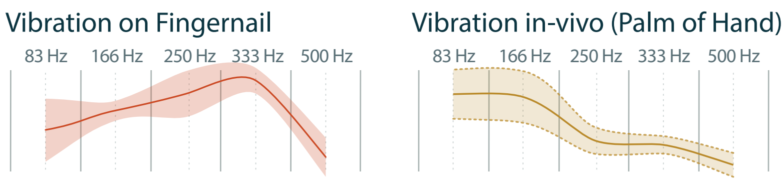

Figure 53: Comparison of vibration applied to fingernail and vibration applied internally via implanted magnet using the *Magnetips* system. Solid lines are mean response. Shaded area is the corresponding 95% confidence interval (left) and two standard deviations (right). Y axis is perceived magnitude.

DURATION Within the ranges investigated (2 ms to 12 ms), duration had a linear positive relation with perceived strength for both on-body (see Chapter 2) and implanted feedback (see Chapter 4). As the pulse duration increased, the haptic signal was perceived more strongly. The difference in duration was not perceived.

For the studies using Magnetips*, we manipulated the frequency of a square pulse-burst. This is subset of ways in which* timbre *can be manipulated. In this context, the terms* timbre *and* frequency *refer to the same general concept.*

TIMBRE Frequency has a non-linear effect on perceived strength. This is observable in the *Magnetips* experiment (Chapter 2) and the evaluation of the implanted magnet (Chapter 4) as well as the *Haptic Textures* data (see Chapter 9, Figure 38). While the *Magnetips* data are roughly what we would expect based on prior psychophysics studies [149, 187, 188], the perception of the in-vivo vibrotactile feedback was markedly different. The experienced vibration on the fingernail was strongest around 333 Hz and grew weaker with higher or lower frequencies. The implanted magnet lead to a strong response to 83 Hz and 166 Hz and was experienced as weaker at higher frequencies (Figure 53, see also Figure 16.).

13.1.2 *Actuation of Proxy Object*

Rather than actuate the skin directly, in most of the work presented in this thesis the skin is actuates by means of a proxy object manipulated by the user.

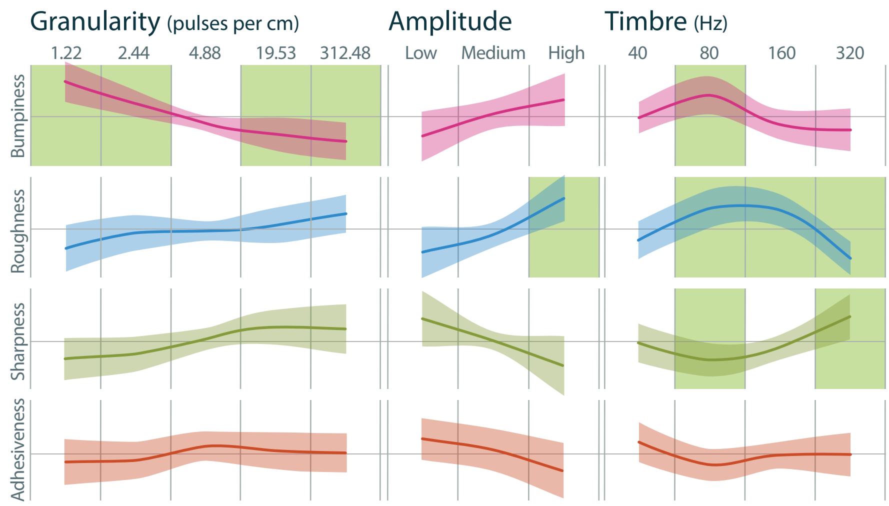

Figure 54: Overview of results from re-analysis of *Haptic Textures* data. Thick lines are mean results, shaded area represents a 95% confidence interval. Feedback parameter levels at which an experience lead to significantly different ratings, compared to the overall average, are highlighted in green.

GRANULARITY granularity describes the number of pulses that a user is subject to relative to a set amount of motion. In the experiments of Chapter 7 we described it in units of pulses per cm (p/cm).

Kildal varied the regularity *in his study of compliance [90], however the analysis does not allow inference on how it effected the experience.*

We assumed granularity would correspond closely to the surface features of a material. Consequently we expected granularity to have a strong effect on experiences and that different types of experience could be differentiated clearly based on their corresponding granularity. We found that this was not the case. While bumpiness was strongly associated with low levels of granularity, we did not find any other clear links between granularity and a specific experience (see Figure 54).

A sub-dimension of granularity is *regularity*. In our experiments, we always had regular signals. That is, the granularity hat a set, constant, value. This was sometimes commented on by participants who felt that it made the experiences feel mechanical or somehow "made by humans", as opposed to something one might find in nature.

AMPLITUDE We assumed that the amplitude would merely impact the strength or salience of the experience, but would not influence the quality of the experience. What we found was that an increase in amplitude lead to the strongest increase in the experience of roughness, followed by bumpiness. Adhesiveness and Sharpness were comparatively barley effected. Roughness was effected significantly more by high amplitude than the other experiences (see Figure 54).

TIMBRE By timbre we refer to the composition of the pulse in the frequency domain. Our initial reason of exploring timbre was that we associated it with the envelope of the pulse, or the pulses shape. We assumed that this would correspond to the shape of a surface feature. We found that timbre was much more versatile than we expected. Timbre had a strong effect on how strongly the signal was experienced, and changing the properties of timbre lead to the strongest differentiation in experience (see Figure 54).

13.2 INPUT MAPPINGS

An example of how input mappings are typically used in HCI are the difference between a pen and a scroll wheel. Using a stylus to write on a pen-enabled display requires no input mappings – there is a direct correspondence between the physical action and its digital representation. The rotary movement of the mouses scroll wheel however is typically mapped to linear motion of a list or text – the physical action is mapped to a different digital representation.

See also the discussion of transfer functions by Zhai [210] for further details on direct and indirect mappings.

When I speak of mapping, I am referring to which aspects of movement are used to create the haptic sensation. I differentiate between *direct* mappings, where there is a one to one correspondence between movement and feedback and *indirect* mappings in which the movement is used to manipulate some hidden variable or transfer function which the haptic feedback is based on [209, 210]. I also differentiate between isometric and isotonic as suggested by Zhai [209, 210] - where isotonic actions require force with contraction of muscles and isometric actions require force without the muscles changing length.

13.2.1 *Direct*

I speak of direct mappings when no transformations are applied to the measure of movement before generating the feedback. The resulting vibrotactile signal directly corresponds to the user's movement.

PRESSURE (ISOMETRIC) The pressure-based device presented in Chapter 12 is the only isometric device I designed. As expected, based on previous literature [90], the resulting experience was of compliance. When stepped on, it felt like walking on a wooden floor that has some give. When pressed, it either felt as if the object was compliant or as if the pressure actuated some internal mechanical mechanism.

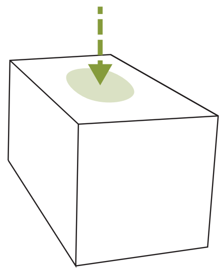

Figure 55: Pressure: The device is rigid and measures the amount of pressure exerted on it. Changes in pressure create haptic impulses.

TRANSLATION (ISOTONIC) I built two devices that use translation as input for generating material experiences. The first was the slider used in *Haptic Textures* (Chapter 7), the second was the pointing device when in translation setting, explored in *Pulse Trains* (Chapter 10).

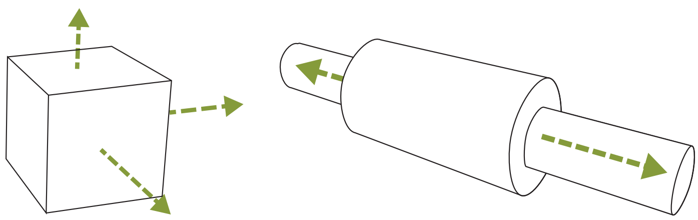

Figure 56: Translation: The device can be moved with minimal or no counterforce in at least one direction. Changes in position create haptic impulses.

The main difference between these two devices was that the slider provided users with normal force, if they pressed perpendicularly to the sliding direction, while the pointing device was completely unconstrained.

Adding haptic feedback to the slider created an experience of friction between the track and the slider. Adding haptic feedback to the non-grounded pointing device provided the experience of pushing through a medium. This is similar to how one experiences textures when writing with a pencil. As the pencil moves over the paper, the pencil is vibrated from its tip. Participants in the *Pulse Train* study often assumed the vibration originated from the opposite end of where the pointer was being held (See Chapter 10 for more details). Participants said that moving the device felt as though it required more force.

ROTATION (ISOTONIC) I created two rotation based devices, the pointing device from the *Pulse Trains* experiment during the rotation condition, and the standalone version reported on in Chapter 12.

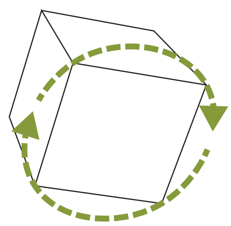

Figure 57: Rotation: The device can be freely rotated along at least one access. Changes in angle create haptic impulses.

These devices behaved essentially the same. Participants typically felt that the motion was somehow mechanically complex, or that the device had traction when being rotated. This was either attributed to the environment – as if the device was rubbing against something invisible – or to the device itself – as if there were internal components set in motion. Typical descriptions included that it felt like a rotary dial or a ratchet. The vibration was often described as emerging from the centre of the device. Again, with the vibrotactile feedback present, participants described that the device felt as though it required additional force to be moved, or as though it was made of heavier-higher quality materials.

BENDING (ELASTIC) *ReFlex* (Chapter 5) could provide haptic impulses relative to the amount it was bent. As the device itself provided a counterforce, this action is not isotonic. Because the device changes its shape when pressure is exerted, it is also not isometric. Zhai calls this type of interaction elastic [209].

With the added haptic impulses, bending the device felt more crisp. The flexible parts of ReFlex either are plastic or feel like plastic. Consequently the bend-action feels very smooth, users only received very few tactile cues when bending it. With added feedback, bending it felt crisper and users received more tactile feedback on their action. Some interpreted this as the device being easier to bend. Others associated

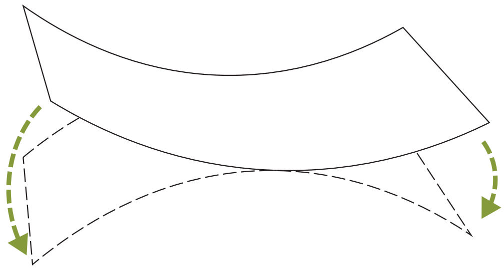

Figure 58: Bending: The device is elastic and can be bent. Changes in the devices shape create haptic impulses.

this experience with a fibrous material, which they felt was harder to bend. Most users expressed that they had greater control over the bending in the presence of the haptic feedback.

13.2.2 *Indirect*

Here the motion is used to generate an interim value. The vibrotactile feedback is based on the interim value or transfer function.

See Section 6.1 for a listing of all feedback and input mappings used.

RATE CONTROL When we provide haptic feedback at a fixed frequency it is experienced as vibration. This was true in the absolute feedback condition, where the frequency of the pulses was determined by how far ReFlex was bent (if the device is flat, no vibration is present, and the more the device is bent, the higher the frequency of the haptic impulses).

On of the ways of using bends as input for controlling *ReFlex* is rate control. When using rate control, the bend-state of the device controls the speed with which the cursor moves - so, when the device is kept in a fixed shape, the curser moves at a fixed speed.

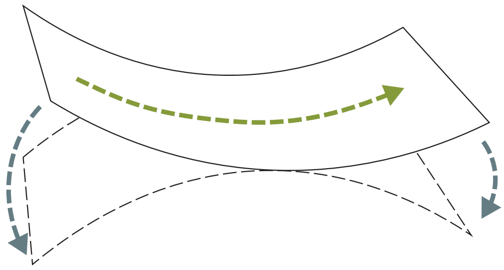

Figure 59: Rate Control: A cursor moves at a speed corresponding to the bend state. The device vibrates at a frequency corresponding to the same bend state. The faster the cursor moves, the more frequent the user experiences haptic pulses.

Combining rate control (bend-state controls speed of cursor) and absolute feedback (bend-state controls frequency of pulses) created a new material experience. It felt as though the cursor encountered friction when moving over the screen. This experience was similar to the slider used in *Haptic Textures*, and I imagine that the results of the *Haptic Textures* study would directly apply to this scenario (Chapter 7).

More details on this mapping can be found in Section 10.4.

PROJECTION The *Pulse Trains* study was set up similarly to the rate control scenario described above. Here the user could move an imaginary point over an imaginary surface, similar to how the light-point of a laser pointer might move over a far-away wall. We imagined the resulting experience might be like touching a far away wall, or as if the imagined pointer was subject to friction.

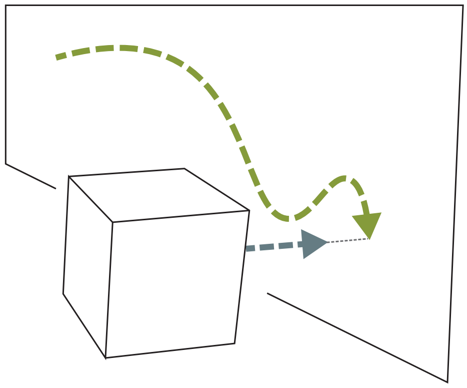

Figure 60: Projection: The device can be freely moved. The intersection point of an imaginary ray emitted from the device and an imaginary surface is calculated. Haptic pulses are created when this intersection is moved.

This did not occur. Participants found it difficult to understand this mapping. Even when we explained it in detail, participants found it difficult to make sense of and perception-shifts usually did not occur.

Multi-Modality for Indirect Mappings The projection condition makes it appear that people have trouble understanding mappings with a non direct transfer function; mappings which use an interim variable. Even if people understand the metaphor of the projection condition, the perception-shift does not occur. This could suggest that a metaphor needs to invoke a motor memory for the shift to occur.

I discuss this point again in Section 14.4.1

On the other hand, the rate control condition was easily understood. This suggests that if the mapping is not understood purely on a motor-level, presenting users with a multi-modal stimulus creates a change in how the interaction is interpreted, which allows the perception-shift to occur.

The Isotonic – Elastic – Isometric Spectrum An interesting observation is that adding motion-coupled vibrotactile feedback to isotonic devices makes people experience them as having some type of counter force, such as friction, resistance or weight. Adding motion-coupled vibrotactile feedback to isometric devices on the other hand, makes these devices feels as though they have less-counterforces. They are experienced as more elastic or compliant than they are. Finally, for the elastic device, the feedback was interpreted either way, depending on user and context. Based on prototyping with various stiffness levels of ReFlex, I believe this to be a coninuous spectrum.

13.2.3 *Experience*

Adding to the perceived material experiences of an object with motion-coupled vibrotactile feedback, transformed not only the object is experienced; it transformed the entire way in which interaction with the object was experienced. Participants often spoke of a greater sense of agency or better control in the presence of the vibrotactile feedback.

CONTROL We asked if the material experiences created by us had any effects on targeting performance in Chapters 6 and in Chapter 10 we asked if the material experience helped them in drawing symbols in the air. For both conditions the participants answered with a clear yes. However in the experiments using *ReFlex* we could find no such measurable effect. I appears as though the changed material experience can provide the user with the experience they would have, if they had more control over the interaction. However, as – mechanically – the interaction is unchanged, this effect appears to be purely subjective and does not translate into improved performance.

AGENCY Some participants pointed out that they receive more information regarding their movements in the presence of the vibrotactile feedback. This contributed to a greater sense of self-perception and agency. Something I also observed with all my prototypes is that they are somehow very satisfying to use. Participants often mentioned that when interacting with them, if the vibrotactile feedback was present, the interaction becomes more enjoyable, and without the vibrotactile feedback, the interaction somehow feels empty. I attribute this additional satisfaction participants describe to the experience of increased agency.

13.3 CONCLUSION

This Chapter presents an overview of input mappings and output parameters explored in this thesis. The results range from very basic low-level observations of parameters that influence how strongly we perceive vibrotactile feedback – I highlight known links between frequency and perceived strength of vibrotactile feedback and point out potential unknown mappings for in-vivo vibration – to exploring how output parameters are linked to qualitatively distinct experiences.

I show how changing properties of the vibrotactile feedback in the frequency domain can create qualitatively distinct experiences. I then show how not only the feedback parameters, but also the input mappings influence the experience. I show that indirect input mappings appear to require another sensory modality to work well and that the isometric-isotonic spectrum has a strong effect on how the feedback is interpreted.

Finally I point out that feedback parameters and input mappings together augment not only the objects a user is interacting with, but provide the user a heightend sense of control and agency. The higher level implications of these results will be discussed in Chapter 14, a fictional system which uses these results directly is presented in 15.

14 THEORETICAL REFLECTIONS

A steadily growing stream within HCI emphasizes the importance of the body and embodied cognition. Around 2005 the term "embodied interaction" became commonplace at CHI and has remained prominent ever since [77]. While one might expect scientific approaches towards HCI to have have testable models (such as Fitts' Law) and clear success criteria (such as reducing task completion time), embodied interaction does not. This has led to a paradoxical situation where embodied interaction has become recognized as an important direction in HCI, but we do not have the tools or knowledge to provide a decisive answer to the following question: *From an embodied interaction perspective, is this interaction designed well?*

It frustrates me how difficult it is to make clear and simple statements about embodiment in HCI. In this section I will outline observations, ideas, and concepts I find important when considering embodiment in HCI. I will use these to outline steps towards an empirical approach to embodiment which I believe could in the future resolve my frustration and be useful in designing, discussing, and evaluating interactions and technologies in the context of our embodied, physical access to the world.

While I sometimes make excursions to discuss other sensory modalities, all my observations are only regarding haptic perception. While many of the observations and ideas presented here might be transferable to other modalities, the limits of such transfer are not clear.

My perspective is intentionally narrow compared to the breadth of work published under the umbrella of embodied interaction. Some sources which have shaped my perspective include *What Things Do*, by Peter-Paul Verbeek [186], the *Phenomenology of Perception* by Maurice Merleau-Ponty [117][1] and *Knowing Hands* by David A. Rosenbaum [147]. However, I do not intend this Chapter to be a comprehensive overview of how to apply these (and other) works to HCI. Instead, I pick and choose the bits which I found useful and ignore the rest. Similarly, I do not intend to deliver a comprehensive theory of embodied interaction, nor a comprehensive overview of the related literature. What this Chapter attempts to do is outline initial steps towards an empirical perspective on embodied interaction.

In this chapter, I will first discuss what happens when we perceive a material property such as hardness, texture, or shape. Based on observations made during my research, I highlight properties of perception which I believe we need to consider when designing embodied interactions.

1 For those who find Merleau-Ponty hard to stomach, I recommend reading the 2012 translation by Donald A. Landes. This newer translation manages to both make the language more accessible and highlight the structure of the arguments, resulting in a much more pleasant reading experience than previous translations.

Following this discussion, I will briefly outline problems I see with current approaches to embodied interaction, before introducing theories and perspectives which I believe we should be considering. Specifically, I provide arguments which explain how – even though all interaction is necessarily embodied – there is utility in discussing the extent to which an interaction is based upon or leverages this embodiment. I then discuss theoretical perspectives on the unit of experience, highlighting that experience is an emergent phenomenon.

I argue that if we wish to design experiences, we need to explicitly design the interactions which constitute our perception of the world. For example, if we wish to design the experience of increased weight, we need not physically changing the weight of the object. I argue that we can take note of the interactions between the body and an object with the target weight, and design our system in such a way that this type of interactivity is supported. The experience of weight will then emerge from the interaction.

I focus on interactions that occur on time-scales shorter than those typically considered within the domain of HCI. This distinguishes my approach to embodied interaction and interaction in general from other perspectives on embodied interaction in HCI.

The thoughts and ideas presented here are not intended to add up to a full-fledged theory yet. I consider what I present here to be a starting point for a dialogue which might in the future culminate in an empirical theory of embodied interaction, but much work is required before we arrive there.

14.1 HOW IS A TEXTURE PERCEIVED?

I have previously focused on parameters for vibrotactile feedback design, and mappings between motion and perception. The experiments of Chapters 7 and 10 also point to a number of general properties of haptic perception. Here, I discuss how the physical world, perception, and the resulting experience relate to each other, and highlight some of the less intuitive observations.

I distinguish between "perception" – the action or process of perceiving – and "experience" – that which emerges from perception. The importance of this distinction will become clear throughout the Chapter.

14.1.1 *Perception is Active*

If a person is asked, "How do you feel the texture of a material?", they might answer, "By touching it". In casual conversation, I often meet strong resistance when I claim that "To understand the texture of a material, one also needs to know what it feels like to move one's finger over it" [170]. While this claim is scientifically not contentious, it contradicts the common conception of experiences: that perceptions are something that - almost inevitably - happen to us. We might cover our eyes to not see something, but if we then allow our eyes to wander, we see it whether we want to or not. There appears to be no conscious effort involved in perceiving.

The observation that we need to actively move our finger over a material to perceive its texture suggests a different model of perception. In this model, perception is not something that merely happens. Rather, it is an action shaped by our intentions. A person does not merely *have* an experience of texture, a person *does* an experience of texture. The idea of active perception is often traced to Merleau-Ponty and is famously demonstrated in images created by Alfred Yarbus [207]. Yarbus shows how the intention that one has when glancing at a picture influences the patterns the eyes trace when looking at it. In the picture of *The visiting stranger* by Ilya Repin (Figure 61, left), the pattern traced by the eyes when asked to asses the social relations between the people in the picture (Figure 61, centre) and their material wealth (Figure 61, right) are clearly distinct.

Figure 61: Left: *The Visiting Stranger* by Ilya Repin. Center: Path eyes trace during free examination. Right: Paths traced by eyes when asked to estimate the material circumstances of the family. Each path shows 25 seconds of eye-movements [207].

These examples are well known, but are typically associated with higher level cognition. Surely, assessing the material wealth of people in a picture is a different process from perceiving a texture? However, I argue that considering such examples of active perception as somehow more complex or higher level is misunderstanding the significance of these images. The viewer is not looking at multiple discrete elements of the image, the viewer is engaged in a continuous interaction with the image. The viewer's perception, be it an impression of the material wealth or an impression of the social relations or anything else, is shaped both by the viewer's intention and the features of the image. The viewer's intention and the image's features together guide how the eyes linger on or move over the image.

The same is true in the perception of textures. The texture experience is not simply there and somehow transported into our consciousness. We need to interact with the material to perceive the texture. This is not merely a philosophically interesting point - it is the reason why the systems presented in this thesis [166, 170, 171] work, and it is the very basis of much of the most exciting work in haptics [39, 90, 134, 146, 206]. These systems, which provide users with material experiences, are not designed in a way that these experiences are ‘played back’ on a static user. Instead, they provide a space which the user can explore while they provide feedback in reaction to the user’s movements.

In conclusion: perception is not something that happens to the body. Perception is an activity that the body engages in. This means that the first step in designing a material experience is deciding which parts of the user’s motion should engage in the perception and how to measure it. As perception is a physical action, we can only start designing material experiences once we have sufficiently precise measures of the users’ body [168].

14.1.2 *Perception is not a reflection of the world*

While we have started correcting the colloquial model by pointing out perception’s active nature, that model contains another implicit fallacy. For example, we might speak of textures as a property of a material. We might also speak of textures as an experience that people have. There is the implicit assumption that perception somehow translates this material property to an experience. The various prototypes I built, and their evaluations, point towards a much stranger model of perception.

When interacting with ReFlex [171] or the Vibrotactile Slider [170], people are presented with vibration. However, people do not experience these vibrations the way they might perceive a vibrating phone. Instead, they report experiencing changes in the material properties of the devices. This should not come as a complete surprise. Research has shown that the vibrations created when interacting with a material are the basis of differentiating between textures [12]. Romano and Kuchenbecker demonstrated how this link between vibration and perception could be leveraged. They recorded the vibrations created by moving a pen over textures and played these vibrations back when a pen was moved over a non-textured surface. Participants in their experiments reported that they experienced textures [146].

Incidentally, this is why we have evolved fingertips. Research has shown that they amplify the vibrations created through interaction to allow us to distinguish better between materials [37, 107].

These observations suggest that we do not have direct access to the texture of a material. What we have access to is created by interaction of the body with the material. The act of perceiving creates the subject of perception. The act of perceiving gives rise to a physical phenomenon - vibration - which is related to the texture but distinct from it.

Again, this is not just an interesting philosophical point. It is the very basis which might in the future allow us to create virtual worlds which we can not only see, but also touch. The observation that our experience of texture is not based on a material, but on the interaction with it, suggests that we can create material virtual worlds without the requirement of recreating the actual material.

Once we have precise measures of the user’s actions [168], we can reason about [170] or model [146] the type of interactivity the user would experience when performing these actions in the presence of a given material. Rather than recreating the object which would cause such a material experience, we can now recreate the interactivity which leads to the experience.

14.1.3 *Experience is not a reflection of perception*

My favorite straw man might now declare, “But textures exist! Surface spectronomy is a scientific field. We know a lot about textures. The colloquial model may be flawed, perception may be active, perceptions might literally be created through this interaction, but at the end, we must perceive an element of something which is truly and objectively out there." The idea that our experiences are somehow equivalent to the objective world is the last implicit assumption often encountered about perception which I argue to be a fallacy, based on my experiences with haptic devices.

Understanding this encoding/decoding process might allow us to create compression algorithms for textile experiences, much like mp3 and similar encoding algorithms compress audio.

On a pragmatic level, there is a disconnect between our experience (a texture) and what is mediated by perception (a vibration). The vibrations that trigger Pacinian cells in the fingertips are somehow translated into a contextualized experience. This suggests that features of the texture are somehow encoded in vibration and then decoded into an experience. While this is indisputable, the malleability and ambiguity of the experience distances it from the perception even further.

I still clearly remember the first time I saw a Necker cube. I do not remember my age, but it was before I started school. The Necker cube looked like an unorganized bunch of lines. I asked my father what they represented and he told me it was a cube. I remember thinking that that made no sense – it’s just jumbled lines. And then, all of a sudden I saw a cube. In my mind I can now change the cube to be viewed from the top or from the bottom. But I cannot un-see the cube. I am unable to experience that lump of jumbled lines. The only thing I have left of them is the memory.

Peter-Paul Verbeek [186] argues that the Necker cube has yet another stable state: that of a six-legged spider in a hexagonal web. While I can intellectually appreciate the point, I don’t see the spider.

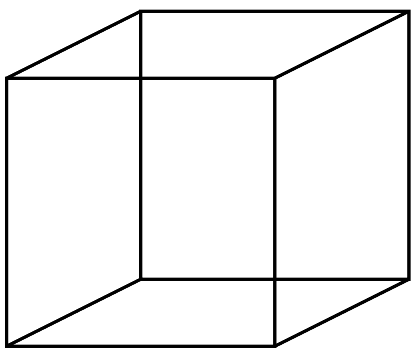

Figure 62: Necker cube

This is similar to how most people interviewed in Chapter 10 described their experience with the haptic pointer. In the "translation" condition of the *Pulse Trains* experiment (Chapter 10), the device vibrated at a speed proportional to its displacement relative to an arbitrary point. When people tried this, often they initially expressed confusion: "Hm, it vibrates when I move it." Then, typically, a perceptive shift occurred: The experience of vibration receded and gave way to a material metaphor "It feels like I am pushing through some medium" (see more detailed examples in Chapter 11). Here we have a clear shift in experience, but the physical material world is unchanged. The act of perception is also unchanged: one perceives the vibrations by moving the device. What has changed is the experience of it.

This suggests that our experiences are organized in higher level units, as they might become relevant to our day to day activities. For example, when sliding a finger over a surface, the specific coupling between finger movement and resulting vibrotactile stimulation of the fingertip is usually not of interest, but understanding the friction of the surface may be of importance. When one recognizes that the vibrotactile feedback follows the same regularity as activities one is familiar with, the vibrotactile feedback is organized accordingly. While the feedback created with the haptic pointer (Chapter 10) has nothing close to the fidelity of a real-world material interaction, people realize its material-ness in the same sense as one recognizes the cube-ness of the Necker cube.

This shift in experience cannot be attributed to empirically measurable features of the world around us. No matter how hard we search, we will not find this change of experience in the world. However, this shift in experience is typically desirable, and is required if one's goal is to create material experiences. It is the difference between the user thinking *Oh, the controller is vibrating, it is providing me with a signal that means that I have reached the edge* and *I can feel the edge using the controller*.

As this shift in perception is not something we directly control, we must instead nudge the participants into performing this shift. To do this systematically, and to find empirical methods of achieving it, I argue that we need theoretical guidance. The rest of this section will outline initial steps towards such a theory.

14.2 TOWARDS AN EMPIRICAL PERSPECTIVE ON EMBODIED PERCEPTION

In this section, I will first outline why current theories and frameworks around embodied interaction are not satisfactory to me. I will then examine some ideas which help me think about embodied interaction. Finally, I outline my own perspective of embodied interaction, as well as qualitative and quantitative design constraints, open questions, and predictions which might be empirically addressed. The aim of this section is to provide a first step towards a tool which might help us to systematically and empirically expand our understanding of how HCI might leverage the embodied nature of perception.

14.2.1 *Problems with Embodied Interaction*

It might appear obvious to look at embodied interaction as a theoretical framework for explaining and leveraging the observations made in Section 14.1. In fact, embodied interaction, as introduced in Paul Dourish's book *Where the Action Is*, [46] uses a phenomenological framing – similar to my own line of reasoning – to highlight the commonalities and create a common framework for tangible and social computing. However, in its current state, I find myself lost in the theory and various interpretations of embodied interaction, as they do not provide me with a scientific framework which might help push my work forward.

Generally speaking I encounter two main problems with embodied interaction: (a) no one really agrees upon what embodied interaction is, and (b) there is currently no scientific theory of embodied interaction.

(a) No one really agrees upon what embodied interaction is

I don't mind there being conflicting theories and perspectives, and I don't mean to suggest we need to eventually arrive at a single canonical theory. However, it does make talking about embodied interaction difficult. If we are speaking of the same concept, we need to find ways of resolving conflicting claims. If we are speaking of different concepts, we need to identify these concepts and understand how they relate.

Van Dijk points out that there are considerable differences and conflicting claims in how embodied cognition is relevant to interaction design [44]. I have encountered the term referring to interactions which require full body engagement, for interaction with tangible objects, or merely referring to the banal fact that part of an interactive system is a physical object. The confusion around the term is mirrored in a literature review of the term *interaction* [77] where the word "embodied" is used both to describe a design approach as well as to describe a property of a system.

The confusion is amplified even further in informal discussion of the topic at conference venues. If a paper compares two different user interfaces and makes claims to these having different degrees of embodiment, this paper might be criticized as missing the point, because "All interaction is embodied", so such comparisons cannot be made. This is rightfully countered by the comment that "If one considers all interaction to be equally embodied, then one should probably not be using embodiment as something one values in design in the first place".

(b) There is currently no scientific theory of embodied interaction

There can be substantial debate as to what a scientific theory is. Without going into great detail, I will use a commonly accepted definition: *A scientific theory is an explanation of an aspect of the natural world that can be repeatedly tested and falsified.* To the best of my knowledge, such a theory for embodied interaction does not exist.

The concept of embodied interaction, as popularized by Dourish's (2001) seminal work [46] was used to demonstrate that there is a commonality underlying tangible and social interaction. However, Dourish does not provide specific guidelines or methods on how these insights might be applied. Since then, there have been a series of works which attempt to distill why the human body is relevant and how this might reflect in design. For example, Klemmer et al. [96] present a collection of themes, ideas, and principles related to embodied interaction. Van Dijk and Hummels [78] present seven design principles for face-to-face, embodied sensemaking technologies. They then present seven principles for designing for embodied *being in the world* in general and distill these into a framework [43]. This framework and its principles have been successfully used as design tools. Van Dijk demonstrates how the principles he suggests might be used as thinking and reflection tools in the design process. He notes how these principles helped him arrive at non-intuitive design choices. In sum, the strength of embodied interaction – as suggested by Dourish and the various principles and frameworks based on embodied interaction and embodied cognition in general – lies in opening up the design space and providing designers with thinking tools (e.g.: [112, 119]).

While Klemmer et al. [96] claim that their principles are intended to guide generation of ideas as well as the evaluation of those ideas, it's not clearly described how that evaluation is supposed to take place. As the principles presented by Klemmer [96] and others [43, 78] are generally not falsifiable, and because there are no clear metrics of what successful embodied design entails – surely traditional metrics such as user satisfaction or task performance provide information orthogonal to embodied design – the main tools remaining for evaluating the success or failure of a design are rhetoric and persuasion.

I am fully aware that persuasion and rhetoric are an important part of the scientific method – in fact, rhetoric and argumentation are, by and large, the bases of this chapter. Conducting Science in a postmodern world, I urge us not to throw out the baby with the bathwater. Recognizing Science as a social construct does not absolve us, as scientists, of rigor in our work and empirical grounding in our theories.

I approach problem (a) by narrowing the scope of embodied interaction. Van Dijk provides a useful first step in resolving the confusion around what embodied interaction is. He identifies three "flavours" of embodied cognition used in HCI [44]: *socially situated practice*, *distributed representations and computation*, and *sensory-motor coupling and enactment*. I suggest that these three are not variations of some shared underlying principle, but rather that they have a hierarchical relation to each other. I see "sensory-motor coupling" as the core principle. It is through such coupling that we gain consciousness of the world. Only once sensor-motor coupling occurs can we access "distributed representations". For a "socially situated practice" to emerge, we need not only understand that there are *objects to* in the external world, but we also need to recognize that there are *others*. I therefore see socially situated practice on the highest level of the hierarchy I am suggesting. As I believe sensory-motor coupling to be the most fundamental element of such a potential grand theory of embodied interaction, I argue that any scientific approach should focus on it first. Once we have greater understanding of what sensory-motor coupling means for interaction design, we can extend it to gradually encompass higher level concepts. Regarding the perceived contradiction between, on the one hand, the view that everything is embodied, and, on the other hand, the perceived loss of utility of embodiment for design, I argue that there is no contradiction. All interaction is embodied. However, the extent to which a particular interactive system explicitly leverages this might vary. I will expand upon this idea in section 14.3.1.

I approach problem (b) by sticking out my neck and suggesting some metrics and conditions for successful sensory-motor coupling and some easily falsifiable claims about the required conditions for sensory-motor coupling to occur. I am not suggesting that these are all relevant metrics and conditions, or even the right ones at all. I do not mean to suggest that they are particularly original or profound. In fact, they are rather mundane and probably obvious to most who have spent any amount of time thinking about the topic. But they're a start. My hope is that with time, some of these will be built upon, while others are rejected; that new ones will be added; and that the relationships between them will be explored on an ongoing basis.

14.3 BUILDING BLOCKS OF A THEORY

Here I present a series of ideas and perspectives on perception and interaction which might be used for explaining the observations made in my experiments.

I discuss qualities of perception to highlight that perceptive acts have multiple, qualitatively distinct aspects. This can help position 'embodied' theories of interaction relative to 'cognitive' theories of interaction without invoking mind-body dualism.

I then discuss units of experience. I highlight the emergent nature of experience and suggest that experiences cannot be reduced to the discrete elements they are constituted of. Experiences emerge from perceptual activity which we are not consciously aware of.

I believe that this perceptual activity should be considered when designing for the body. To highlight how this approach differs from most other approaches in HCI, I discuss temporal bands of human activity. I point out that there is utility in considering the intersection of the biological and cognitive temporal bands. I refer to this as the "perceptive band".

14.3.1 *Qualities of Perception*

A number of theories distinguish among various aspects of perception or action. I present a somewhat arbitrary selection of some of these ideas which resonate with me and highlight a distinction I believe we need to make when discussing embodied interaction.

I discuss qualities of perception to explain what it is that I am interested in.

MICRO- VS MACROPERCEPTION Don Ihde speaks of micro and macroperception as two distinct, but closely related and intertwined, dimensions of experience [186]. He suggests that *microperception* is the bodily dimension of sensory perception, while *macroperception* consists of the frameworks within which these bodily perceptions are made interpreted. While these two types of perception can be distinguished from each other, they cannot be separated. There can be no bodily perception without it being interpreted, just as there can be no interpretation without the bodily perception (see Verbeek [186] pp. 122-123).

EMBODIED VS HERMENEUTIC MEDIATION Ihde draws another, similar distinction when he discusses how information might be mediated by technology: he suggests that there are two types of technological mediation, embodied and hermeneutic.

Ihde uses the example of a dental probe to describe *embodied mediation*. When a dentist uses a dental probe to detect cavities in a patient's teeth, the dental probe extends the sensitivity of touch, so the dentist feels the cavities by means of the probe.

As an example of *hermeneutic mediation*, Ihde uses the example of a thermometer. The thermometer does not enable us to feel the temperature, but the thermometer represents the temperature to us in symbols which we interpret.

In embodied mediation, our senses are extended to perceive what we are interested in. In hermeneutic mediation, the world around us is modified so that the information is presented to us (see Verbeek [186] pp. 123 - 126).

PROXIMAL VS DISTAL INTERACTION A similar distinction is made by Rosenbaum [147] when discussing ways in which sensory stimuli and their resulting actions are linked. He suggests that actions are prompted by stimuli in various ways. He distinguishes between proximal (direct) and distal (indirect) triggers of actions. A *proximal interaction* with a water bottle is seeing the bottle and instantly appreciating that it makes drinking possible. A *distal interaction* with that bottle might be inspecting its label and learning, through inference, that the content can be drunk. Another proximal interaction described by Rosenbaum is the act of pulling one's hand away after touching a hot surface; another distal interaction is that thinking about how Alice's changing size in *Alice in Wonderland* inspired Rosenbaum to write about the short-term and long-term contributions of vision to manual interaction.

What all of these distinctions have in common is that they attempt to differentiate between phenomena which are, for lack of a better word, *close* to the body from similar phenomena which appear to require a more cognitive explanation. What they also have in common – something which is also present in my own argumentation – is a certain clumsiness as they attempt to draw these distinctions while concurrently trying to avoid evoking a mind-body dualism.

Without resolving this conflict, I believe it is still safe to say that interactions with the world have different qualities, some leverage our embodied existence more (proximal, embodied, microperceptions), and others less (distal, hermeneutic, macroperceptions).

14.3.2 *Units of Experience*

I discuss units of experience, as these ideas may help describe what occurs when vibration is coupled to motion; also, they might help us evaluate the experiences created by such systems.

Something happens when people interact with haptic systems that are coupled to motion. I am unable to provide an objective account of this phenomenon. Instead, I have quoted various first person narrations (Chapter 11). There are a number of ways in which these accounts might be explained. In this section I present three of them.

GESTALT THEORY Gestalt psychology is typically traced back to a 1912 paper by Max Wertheimer, and has remained relevant in modern psychology to this day [190]. It is often invoked in the context of graphical design, with a focus on *emergence, reification*, and *multi-stability*. In this context, emergence means that we first identify the whole, before we might focus on its parts. Reification refers to the generative part of perception: our ability to see and identify objects for which only partial visual information is provided (see figure 63). Finally, multi-stability refers to the property of perception that – if presented with an ambiguous stimulus – we tend not to perceive the ambiguity, but instead our perception switches between the various stable states. A well-known example of this is the Necker cube, which I already mentioned earlier (see Figure 62). It should be noted that when gestalt principles are discussed in modern psychology, they are typically referencing the underlying rules and regularities which lead to the phenomena just described [180].

While most examples of gestalt psychology refer to visual perception, related phenomena are common in other perceptive domains. A well known example from auditory perception is that when listening to a sound, a harmonic series, we can hear the fundemantal frequency even if it is not present. Bregman has demonstrated that gestalt principles also apply to auditory perception [22].

Svanæs [175] argues that gestalt principles apply to interactive systems. Using a simple interactive system consisting of one, two, or three squares which switch color from black to white or white to black, respectively, when one of them is clicked, Svanæs argues that users do not perceive the interaction as individual atoms of actions and reactions. When describing the individual conditions, users use comparisons to other interactive experiences, rather than bottom-up analysis of the behavior of the system. Svanaes claims that this suggests users perceive the complete interactive behaviors. He concludes that the experience of

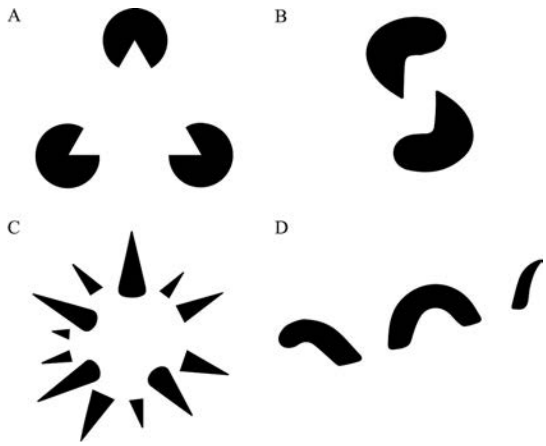

Figure 63: A triangle is perceived in figure A, though no triangle is there. In figures B and D the eye recognizes disparate shapes as a single shape, in C a complete three-dimensional shape is seen, though it is not drawn (By Slehar at English Wikipedia - https://commons.wikimedia.org/w/index.php?curid=2428262).

interactivity has gestalt properties. The experience is not constructed of discrete input behaviors and corresponding outputs. Rather, the smallest unit of experience is a complete behavior, as exemplified by the toggle behavior of a light switch.

TOP-DOWN PROCESSING Rosenbaum [147] discusses top-down processing, which he finds has many parallels to what he calls "embodiment effects." He describes top-down processing as "effects of high-level knowledge on lower level data." A well known example is that *yu cn probly red ths txt eevn tohugh it is grabeld*. The general account of how top-down cognition works is that "higher level nodes receive excitatory input from lower level nodes and repay the favor, so to speak, by sending excitatory input back to their lower level 'friend' nodes." So, for example, the node for the word *this* might be actuated even though the *i* is missing. The *this* node might then send signals back to the *t*, *h*, *i*, and *s* nodes. As a result, a second phenomenon occurs. The *i* in *this* can be recognized more easily than if it occured out of context. This phenomenon exists not only in text recognition, but is found over all areas of perception.

Recent work in neuroscience demonstrates that bistable perception of low level stimuli works according to a similar mechanism. Schneider et al. [153] demonstrated that a stimulus with unambiguous horizontal or vertical motion registered in two distinct columnar clusters of the cortex. When participants were presented with an ambiguous stimulus, the neural activity of the perceived direction of motion matched that of the corresponding unambiguous stimulus.

This supports the idea of top-down processing. When presented with ambiguous or incomplete information, we match it with patterns which we are already familiar with.

PERCEPTION AS A CLOSED LOOP CONVERGENCE PROCESS If we consider perception as an open loop process, if we believe that sensory information flows unidirectionally from our senses to our brain, then there are several phenomena that require explaining. For example, perception takes more time than can be easily explained by a unidirectional information stream. Also, perception is typically accompanied by motion, which – if one assumes perception to be an open loop, unidirectional process – would require that the brain somehow correct for movement, as, for example, the speed at which I move my finger over a material changes the physiological stimulation of my finger, but the quality of the experienced material is unaltered.

"Closed loop" describes systems in which every signal eventually affects its source; "open loop" describes systems in which signals cannot affect their sources [1].

If, on the other hand, we assume that perception is a closed loop system, the longer duration of perceptive acts is easy to explain and movement no longer becomes a source of error which needs to be filtered out, but becomes a feature of perception. Ahissar and Assa [1] argue that a more accurate view of perception is that "the brain triggers the movement of the sense organs, and thereby alters the sensory information that these organs receive. This information is relayed to the brain, triggering further movement of the sense organs and causing the cycle to repeat. Perception is therefore a “closed loop”: information flows between the environment, sense organs and brain in a continuous loop with no clear beginning or end."

They argue that these loops, involving motor, sensory, and neuronal components, are perturbed by new sensory information. As a new feature is introduced, the closed loop system converges towards a steady state around the feature. As the steady state is approximated, the experience arises. Perception is then the process of moving from perturbation to the steady state. Ahissar and Assa [1] mainly use examples from research into haptic perception, with a special devotion towards the whisking movements performed by cats or mice. They argue, however, that their views on perception should hold for other sensory channels as well.

14.3.3 *Temporal Dimensions of Interaction*

I discuss Newell's bands of activity, as they help position my area of interest relative to other approaches in HCI.

To better explain how my view on embodied interaction is different from, and relates to, the more general concept of HCI, it is useful to look at the temporal structure of interactive systems.

Newell [124] suggests that human activity happens on multiple temporal bands, spanning orders of magnitudes: milliseconds for biological activity, seconds for cognitive activity, minutes for rational activity, and weeks for social activity (see table 10).

MacKenzie [108] uses this model to highlight how multidisciplinary HCI is. MacKenzie claims that HCI occurs on all bands, from cognitive to social, and implies that HCI may even be relevant on an even higher scale, the historical one. In line with early HCI theory – such as the *model human processor* by Card et al [28] – MacKenzie [108] completely ignores what Newell calls the biological band, as if it were naturally out of scope.

It takes ∼100 ms for sensory driven neural activity to reach relevant cortical areas. After ∼150 ms one is able to make initial crude statements about a stimulus [202]. The perceptual activity required for this to occur, must happen at faster time-scales.

Months Weeks Days		Social Band
Hours 10 Minutes Minutes	Tasks	Rational Band
10 Seconds 1 Second 100 ms	Unit Task Operations Deliberate Act	Cognitive Band
10 ms 1 ms 100 µs	Neural Circuit Neuron Organelle	Biological Band

Table 10: Newell suggests that human behavior occurs over various magnitudes of time-scales. He organizes these time scales into 4 bands.

If we are to take embodiment seriously, we cannot ignore this biological band. Without it, all other bands become meaningless, as this is the timescale at which our body perceives the world. I suggest that it might be a useful thinking tool to insert another band between the cognitive and biological: a perceptive band, the locus of the dance of stimulus and action through which we perceive the world around us.

In the following sections, I will highlight the relevance of the perceptive band to HCI and how it can be leveraged.

I do not suggest that there is no HCI research that engages with the perceptive band. In some domains, such as eye-tracking, it is quite common (e.g.: [189]). My contribution here is - at best - making the distinction between cognitive and perceptive explicit.

14.3.4 *Connecting the Building Blocks*

How can we leverage our embodied nature? How can we create designs which emphasize the proximal over the distal, the microperception over the macroperception, the embodied over the hermeneutic? Acknowledging that perception is not constructed of discrete actions and reactions, but a closed loop activity in which experiences emerge, how can we make use of this? How might we extend interaction design not only towards the historic timescale, but towards the timescale of the neural network and the biological band?

A first step towards finding out how these ideas relate might be a zooming in, taking a more detailed look at what we consider relevant in designing interactive systems.

In traditional HCI we typically design around concept such as input and output, or user needs. Taking the example of moving a token on the reacTable [86], we might describe the interaction like this:

A user wishes to change the pitch of a tone. The user moves the token to a new location (input) and the tone changes its pitch (output).

These interactions occur in the *cognitive band*. The psychophysics literature suggests that concurrently there are a plethora of other interactions occurring. These other interactions are necessary for the user to successfully move the token from one place to the other; they are the act of haptic perception. These perceptions allow material experiences to emerge:

Initially the hand must establish grasp-contact. The fingertips are typically relaxed. As the fingers and hand touch the token, the fingers conform around the object. The resulting tactile information provides an experience of the token's shape and enables the user to grasp it. When grasping the object, normal force is applied, relative to the expected tangential force which will be present when lifting the token [84]. These forces are adjusted relative to the token's compliance and texture. Changing the force and observing how the change in force influences the change in compression of our fingertips provides an experience of compliance [15]. Micro-movements over the surface of the object create an experience of texture [12]. Finally, the object is lifted, with the weight now distorting the fingertips tangentially. This distortion is compared to the expected tangential force while the grip is continuously readjusting so as to not drop the object or provide excessive force, creating an experience of weight [84].

Here, interactions occur at the bottom end of the cognitive band and start moving into the biological band. The experience of the shape, compliance, friction, and weight of the object occur through interactions which we do not consciously attend to. Yet we are continuously performing them as we complete the task at hand. I believe that interactions at this perceptive band – located between the cognitive and biological – are where we should be focusing our attention, if we wish to leverage embodiment.

There is no value judgement intended. In using the term "embodied" in this context, I make no claim over the value of the research or the quality of the design. "Embodied" merely describes the perspective taken in analyzing an interaction or the emphasis chosen in a design.

I argue that even though all interaction is embodied, not all interactive systems are designed to specifically leverage the physical, embodied, active nature of perception. When I push a button, this is an embodied interaction. As I push the button, the interaction between the yield of the button and the compliance of my skin provides me with an embodied understanding of its force-displacement profile. However, work that frames pressing a button as an open loop system is disregarding this embodied dimension of interaction [133]. If, on the other hand, the button is designed so as to communicate supplemental information in its press dynamics [132, 179], if the design takes advantage of the closed loop perceptive processes [1] which occur when pressing a button, then the design leverages the embodied nature of interaction [106].

14.4 BOUNDARIES, QUESTIONS, AND PREDICTIONS

Here I present conditions which I believe need to be met for a perception-shift to occur. I then point out that there may be a measurable boundary between the perceptive and the cognitive band, which could be experimentally established. Finally I provide some questions and statements which might be addressed by a theory of embodied perception. This section is more vague than I would like it to be. I hope that, as work towards a theory progresses, simpler and clearer statements can be made.

14.4.1 *Required Conditions*

Based on my observation I suggest this set of conditions is required for the *perception-shift* to occur.

CONSISTENCY OVER TIME Users need time to learn and understand mappings. Many of the participants in the *Pulse Trains* experiment (Chapter 10) needed to identify the mapping before their perception-shift occurred.

CONGRUENCE BETWEEN MAPPING AND PERCEPTUAL CHANNEL Our hearing is accustomed to perceiving frequencies which are predominantly shaped by the source of vibration, rather than movements of the body. Classifying materials by frequency would be a reasonable mapping for audio perception. However, in our *Pulse Trains* experiment, constant vibration was mainly perceived as irritating. It is reasonable to assume that we are most skilled in interpreting stimuli when they are presented using mappings we are accustomed to in our day-to-day lives.

MOTOR FAMILIARITY Participants who are familiar to the mapping from their day to day interactions were able to perform the perception-shift fast. While the rotation condition was not obvious to all participants, an avid cyclist immediately explained that it reminded her of the ball-bearing of a bicycle, while a tinkerer suggested that it felt like rotating a potentiometer or a ratchet. Providing participants who did not have such experience with these metaphors did not appear to have the same effect.

CONTEXTUALIZING THE MAPPING People who did not experience the perception-shift immediately could often do so if provided with some mental imagery, such as, "Imagine rotating a combination lock". However, this only worked if participants could use this mental imagery to remember an experience they were familiar with.

As designers, we can only provide part of the conditions required for a material experience to emerge. While we can ensure consistency over time, and chose an appropriate mapping, the user brings their own experiences and background, which influence how the stimulus is perceived. Still, there appears to be some malleability, in the sense that people are able to learn new experiences and integrate them into their perceptive repertoire. Multiple participants reported a tension between things feeling new and exciting, unlike anything they had perceived before, and, at the same time, feeling familiar and natural. A potential avenue to explore are multi-modal stimuli, where familiarity in one sensory modality helps contextualize a novel stimulus in the other sensory modality. A theory of embodied perception should be able to provide a clearer explanation on the roll of motor familiarity, and under what conditions contextualization is able to overcome missing motor familiarity.

Boundaries

The major quantifiable difference between interaction on the cognitive and the perceptive bands might be the time scale. The cognitive band and above are what we typically consider in HCI. It is the level at which people pragmatically get things done. The perceptive band is where we can explicitly design the qualities of the experiences people have while performing a task in the cognitive band.

Newell [124] suggests that the deliberate act is located at the lower boundary of the cognitive band (100 ms). Newell also suggests that the biological band starts at 10 ms. As I suggest above, it may be useful to consider an intermediate band where the cognitive and biological meet, the perceptive band.

If we wish to design for the perceptive band, we need to understand the measurement and update rates we need to reach so that the interaction is merged into a single experience. Indicators that we can look at for guidance include the micro-interactions performed between touching and lifting an object. These typically happen in under 250 ms [84]. However, this fact provides no indicator of the temporal structure of the perceptive processes which occur within those 250 ms.

In practice, systems which attempt to create such interactions report maximum latencies of 25 ms [166] to 60 ms [90], with sampling rates ranging from 125 Hz [170] to 2000 Hz [171] and update rates ranging from 125 Hz [170] to 200 Hz [171]. Due to the *just good enough* attitude in HCI prototyping, these values should be considered as an indicator of the lower boundary of required fidelity. While typically rarely mentioned, the role played by jitter (in this case, the variability of latency) is also relevant [196].

Another area at which we might look for guidance is the design and evaluation of novel musical instruments. For example, Maki et al [111] report that musicians can notice latency of 20 ms or higher when playing a theremin. However, when tactile feedback is present, this value has been reported as low as 10 ms. The effect of latency is however context dependent: musicians are able to perform on mechanical church organs with several hundred milliseconds of latency.

It appears that the exact boundaries might be fluid, and users might associate artifacts created by latency or jitter to material properties. For example, a user in the *Pulse-Trains* study identified the latency, but referred to it as inertia. A theory of embodied perception should not only indicate the boundaries to the perceptive band, but also explain and predict how approaching or crossing these boundaries effects the resulting experience.

Open Questions, falsifiable Statements

BODY ILLUSIONS If someone moves a stick in a continuous motion over a regularly spaced grating, they experience a haptic impulse at a regular interval. If this interval suddenly increases, this would either mean that the movement speed has become higher, or that the grating has become finer. All my experiments were set up so that participants had reason to believe that their body would be unchanged. Any unexpected interaction was, therefore, attributed to the object.

"You're definitely not going to grow another head out of your elbow." "Well, OK, proceed, then." – Nathan Horrowitz

Assuming participants have reason to believe that their body is fixed, and, instead, the world might be changing, a change in material experience might no longer be attributed to changing material properties, but rather to their own bodies. If this is the case, then we might be able to create strong proprioceptive illusions using vibrotactile feedback coupled to human motion. This idea also appears reasonable, considering that our haptic and proprioceptive senses leverage many of the same types of sensory receptors (See also chapter 3). More research into this might be incredibly relevant for virtual reality. Rather than discrete redirected pointing, leveraging visual dominance, one might be able to do continuous eyes-free redirected movement.

STABILITY OF PERCEPTION SHIFT As described in Chapter 11, the perception-shifts were not gradual – there was no fading in or out of the shifted experiences. They either occurred or did not occur, just as one experiences the Necker cube (Figure 62) in either one state or the other, but does not experience gradual transitions between the states. However, there still seem to be degrees in the strength of such emergent experiences. As I described, I have a clear memory of the Necker cube pattern before I could recognize it as a cube. I am unable to return to that previous state. This is different from multi-stability, which one can change at will. It is again altogether different from grbld wirttng, wcih can be regnizcoed as long as one does not pay it too much attention (see also the discussion of *Top Down Processing* in Section 14.3.2).

Understanding this better will be valuable for evaluation. I expand upon this Idea in the example application of Chapter 15

Understanding the intensity with which a perception-shift occurs might provide tools in evaluating the nature of the mediation. Intuitively, I suspect that there is not a continuous gradient of perception-shift strength, but rather a number of qualitatively distinct levels. This too requires further research.

MOTOR MEMORY OF METAPHOR? I have previously argued that motor memories support establishing the perception-shift. A different interpretation of the observations made in the *Pulse Trains* experiment [166] – Chapter 10 – is that the metaphor alone helped create the perception-shift. In practice, some participants used abstract metaphors (e.g.: more/less natural), some tried to explain the experience by approximating the mechanical context which might create it (e.g.: a stick pushing against a rubber ball), and some referred to previous experiences (e.g.: turning the rotary dial on their 3D printer). Understanding the role of motor memories versus metaphors might help better understand the type of sensations which might be mediated through a given sensory channel.

An indicator that metaphor alone is not sufficiant is that users struggled with the projection condition, even once they were provided with the corresponding metaphor. See also discussion in Section 13.2.2.

MULTI-MODALITY The projection condition of the *Pulse Trains* experiment did not work well, even after we provided the users with a metaphor. Providing an example to explain the mapping usually did not help them fully make sense of the perception. In the rate-control setting of ReFlex, however, which uses a similarly indirect mapping, participants did experience a perception-shift. A possible explanation of this is that there was a direct visual mapping, supplementing the indirect motor mapping.

see also the discussion of direct and indirect mappings in Chapter 13.

As I promised falsifiable statements, let me rephrase these *open questions* as statements:

- Once a user is accustomed to a vibrotactile signal relative to their movements, a change in the mapping will be explained either as a change in the material experience or as a change in how the body is being moved. By providing contextual cues indicating that the environment is fixed, the change will be located in the body.

- Perception shifts occur in different levels or strengths. Just as other perceptive effects such as multi-stable images are not continuous, these levels are qualitatively distinct.
- Perception shifts require a stimulus that is designed in such a way that it either matches or can be related to a previous motor experience.
- Perception shifts require a stimulus that is the result of a direct mapping of user action. If such a mapping is not used, it can be supplemented by a direct mapping using a different sensory modality.

14.5 LIMITATIONS AND SCOPE

14.5.1 *This is not a theory*

According to Whetten, [195] the building blocks of theory development are *What*, *How*, and *Why*. In the present case, *What* are the variables, constructs and concepts used. They should be both comprehensive and parsimonious - including all which is relevant while staying as simple as possible. The *How* then addresses the way in which these concepts relate. The *Why* grapples with the underlying dynamics that justify the selection of concepts and relations. Additionally, Whetten suggests that one address the *Who*, *Where*, and *When*. These serve to explain the limits within which the theory applies. A good theory should provide a plausible, cogent explanation for why we should expect to observe certain relationships or patterns.

What I have presented here is an enumeration of concepts (*Whats*), some suggestions on how they might related (*Hows*), and various arguments on *Why* these concepts and relations are relevant. The concepts presented here are probably not comprehensive, and certainly not parsimonious. The relations between them have not been systematically explored. Because of this, the explanatory value of the ideas I am suggesting is still unclear; this is reflected in my collection of open questions and falsifiable statements, which still seems somewhat ad hoc.

However, I do not mean to present a complete theory here. I believe this requires dialogue and reflection, and is simply beyond the scope of this thesis. I intend the ideas collected here to serve as a starting point for such dialogue and as an anchor for future discussion and work towards theoretical contribution to the field of HCI.

14.5.2 *Other sensory channels*

The basis of a large portion of my argument is that we cannot perceive material properties without motion. This aligns nicely with the often-repeated argument advanced by Merleau-Ponty that all perception is active. However, this – very literal – perspective of active perception does not intuitively transfer well to other sensory channels. Surely we hear or see without requiring motion of the body?

While I only mean to make claims about haptic perception, there is an argument to be made that motion is similarly important for other sensory channels, but the connection might be more complex and subtle. For example, the importance of motion in vision goes beyond the patterns that we trace with our eyes when we observe a scene [207]. A study by Stevens et al. injected healthy participants with agents which inhibited eye movements, and subsequently the lead author subjected himself to full paralysis [161]. Stevens reports that the perceptions with inhibited eye movements were "striking and often confusing." Participants had difficulty describing what they saw and often contradicted themselves. Only after about 15 minutes could they report systematically on their observations. These included the visual world disappearing or jerking. In the words of a participant, "The world did not *move* ... it was not as if you had taken the stimulus and *moved* it across the screen. When I moved my eyes up, the whole screen was *displaced* up ... [the stimulus] disappeared and then popped up again in another place." Stevens continues to describe that during the full paralysis condition, "image fading became a real problem". He states that it was only due to inadvertent movements associated with the artificial respiration that the images never faded for long periods of time.

(emphases mine)

These observations are especially interesting in light of a case study of a woman who could not move her eyes [63]. As a result of a congenital condition, she has had not eye movements since birth. Unlike the participants in the experiment by Stevens [161], she was reported to have normal vision. However, she had adapted unusual head movement patterns. The patterns matched the eye movements one would otherwise expect in all but speed.

These observations indicate that the very literal interpretation of active perception may transfer to other senses to a greater degree than one might intuitively assume. Additional research on this matter is required.

This section is largely inspired by Tailor Carman's discussion of how Husserl and Merleau-Ponty view the body [29].

14.5.3 *If everything is embodied, how can some things be more embodied than other things?*

I have somewhat sidestepped the issue of why it is useful to invoke embodiment if everything is embodied. My argument – that instead of the degree of embodiment, we might speak of the degree to which a system leverages our embodiment – still implicitly assumes that there can be a gradient of embodiment. A potential way of explaining such a gradient of embodiment is observing how we, subjectively, have sensory access to our embodiment. While all our actions and perceptions are embodied, some perceptive acts provide us with a higher level of awareness of our embodiment than others. I will argue this by comparing vision to touch. Touch has a number of interesting properties:

- We experience ourselves experiencing touch (*I felt myself touching something cold and wet*) and locate touch on our body (*Something touched my left hand*).
- Our experience of touch leaves traces in the world (*It fell apart when I touched it*).
- We experience touch as both transitive (*The boy touched the ball*) and intransitive (*The boy and the ball touched*), as well as both active (*I am touching the ball*) and passive (*I am being touched by the ball*).

These properties of touch establish how we relate to the world. By experiencing the perception of touch and locating it on our body, we understand that it is our body that is experiencing the stimulus. By spatially locating the perception in our body, we learn that our body has a dimensionality. Because our experience leaves traces, we learn not only that we exist, but also that we have agency in the world. Finally, as we experience touch as something that our body can do, something that is done to our body, and something incidental to the state of our body, we learn that there is a plurality of ways that our body can relate to the world.

In short, touch teaches us that we exist, that we have agency in the world, and that the world and the body can interact. This is in stark contrast to, say, vision, which typically:

- Does not reveal its experience (*I can look at my eyes in the mirror, but I do not see the process of seeing*).
- Leaves no trace in the environment (*Look, but don't touch!*).
- Vision is exclusively transitive and active (*The boy looked at the ball*).

Of course we might also say the ball is being looked at by the boy*, but never* the boys is looked at by the ball.

If the only sensory access to the world were vision, would we know of our own embodied existence? The act of seeing provides no introspection into the process of seeing and therefore provides little information that it is a body doing the seeing. While we perceive the world, we do not learn, through vision, that the world can be manipulated. We learn of only a single mode of engaging with the world.

While embodiment is a precondition for both touch and vision, these senses shape our experienced embodiment in different ways. This suggests that, even if we take embodiment seriously, there may be grounds for speaking of degrees of embodiment. This idea possibly contradicts the previous comments on other sensory channels (Section 14.5.2), and requires further exploration.

14.5.4 *But what about X?*

The perspective I take is narrower than that taken by Dourish [46], van Dijk, [43, 44], Hummels, [78], Svanæs [175] or Klemmer [96]. I believe that there is a trade-off between having a theory which is broad enough to be inclusive to all kinds of work and a theory which is narrow enough to develop testable and falsifiable statements. Further, a theory must be strong enough so that successive work can build upon previous work in depth, rather than simply expanding the breadth of a field. It is my hope that others also explore their first principles and approach the topic of embodied interaction empirically from their own perspectives, or that – eventually – the starting place I am suggesting might be extended and refined so that it gradually becomes inclusive enough to accommodate broader research perspectives.

14.6 CONCLUSION

In this chapter, I initially analyze what happens when we perceive a material property such as hardness, texture, or shape. Starting at first principles, based on observations made during my research, I argue that (a) perception is an activity that we engage in; (b) we do not directly perceive the world, but what we perceive is the result of our interaction with the world; and (c) experience emerges from this perceptive activity, and cannot be attributed to features of the world, or to the act of perception alone.

I argue that this means that there is a bodily physicality to our interactions which deserves specific attention. While I believe that the physical and the cognitive are inseparable, I argue that interactions can have qualities which relate to one more than the other. As examples, I refer to Don Ihde's distinctions between micro- and macroperception and between embodied mediation and hermeneutic mediation, and David Rosenbaum's distinction between proximal and distal interactions.

While I argue that perception is the interaction our body performs with the environment, the interaction is not part of the resulting experience. The resulting experience is always whole, similar to how gestalt psychologists describe the perception of visual shapes. This might be explained by top-down processing effects – that low level features of perception triggers high-level experiences. Finally, I discuss the concept of closed-loop perception, which provides a detailed low-level account of the active nature of perceptive acts and how they relate to experience.

I argue that these closed-loop perceptive acts can form the basis of a theory of embodied interaction. Consequently, when designing or analyzing interaction through the lens of embodiment, we should attend to the temporal patterns of such interaction. Using Newell's bands of human activity, I suggest that there is utility in adding a perceptive band between the cognitive and biological, and that we can leverage this perceptive band of human activity.

While these considerations do not yet make up a theory of interaction, in the light of them I make the following predictions: Because of the reciprocal nature of touching, I believe that by designing micro-interactions accordingly, we can not only shape the perception of the world, but also the self-perception of the body. I also argue that a thorough investigation will find that there are different levels to which novel experiences can be triggered (what I call *perception-shift*). I believe that these are finite and discrete, similarly to how our experience of multi-stable stimuli is always restricted to one of the possible states. I believe that an understanding of different types of perceptive shifts might form the basis of a useful metric for evaluation. Finally, based on observation, and consistent with the physicality of the phenomena I discuss, I believe that for novel experiences to emerge, these must be directly linked to a motor memory. I believe that abstract metaphors will not trigger perception-shifts if the mapping does not match what one is familiar with. However, I believe that multi-modality might allow us to become familiar with new mappings.

See also the discussion of indirect mappings in the previous section

15 PRACTICAL IMPLICATIONS

In this section I explain in general terms how one might implement a system with motion-coupled vibrotactile feedback. This is based on my experience in building the systems presented in this thesis, and in light of the theoretical considerations I shared in Chapter 14. I explain how one might approach such a project, highlighting how the theoretical considerations help in thinking about the design and how they might in the future provide tools for evaluating such systems.

15.1 LEVELS OF INTERACTION: PERCEPTIVE TO COGNITIVE

On a fundamental level, we need to recognize that we engage in all interactions in multiple levels and time-frames concurrently. For my purpose we can simplify this and distinguish between a micro level of interaction – which occurs in the *perceptive band* and a macro level, which occurs in the *cognitive band*. These two levels, the perceptive and the cognitive, are different aspects of the same thing. I am not suggesting that instead of designing for the cognitive level we should now design for the perceptive level. Rather, both are always present - it is merely a matter of our level of analysis which might highlight one over the other. The cognitive level is the pragmatic level of *getting the task done*. The perceptive level is concurrent and is evoked when I ask *what was it like?*

There are of course many more levels of analysis, and depending on ones interest one might go deeper to a cell based level for direct nerve interfaces or higher and consider social or historical bands, see also Section 14.3.3.

A nice example of this can be found in two concurrent papers by Oulasvirta et al. [106, 133]. In *The Neuromechanics of a Button Press* Oulasvirta et al. [133] model the pressing behavior as an open loop process. In doing so they place their work firmly in the *cognitive band*. However, surely we learn to adapt our typing behavior to different keyboards. To do so also requires concurrent closed-loop behaviors, otherwise we would not be able to learn the button's dynamics. Oulasvirta et al. are aware of this, in fact the paper *One Button to Rule them All* demonstrates how one might might approach this in the design of a button with dynamic force-displacement curves [106]. In this paper they engage with the activity of a button press in the *perceptive band*.

Awareness of this distinction helps us make explicit design choices for both domains and can be a useful thinking tool. Knowledge of the temporal constraints can point us towards towards selecting the correct components for the job. Understanding the preconditions for the perception switch, and the perception switch itself, supports the design, and may in the future help in evaluating how the system performs in the perceptive band.

Figure 64: Tidmarsh

15.2 CASE STUDY: AN IMAGINED ADDITION TO THE TIDMARSH LIVING OBSERVATORY

The thesis by Gershon Dublon – Sensor(y) Landscapes: technologies for new perceptual sensibilities *– shares similar motivations to my own work. It is well worth a read [47]*

Over 100 years ago, Tidmarsh – a former marsh in Massachusetts – was turned into a cranberry farm. Today it gradually is being restored to its original state as a wetland. Typically funding for such restoration processes are scarce so there is usually none left over for collecting data regarding the process or the success of such an operation. The Living Observatory project of Tidmarsh is an attempt to do things differently. Researchers of the MIT Media Lab's Responsive Environments group have set up sensor-nodes throughout the area to document and share how the environment is being transformed [114]. The data collected has also been used for creating virtual worlds [66] and, most relevant to my own interests, served as the basis for creating sensory augmentation technologies [47].

I had the pleasure of visiting Tidmarsh in the fall of 2017. I will use this setting to describe an imagined technology which I would love to develop. I will explain how the distinction between *perceptive band* and *cognitive band* helps conceptualize and understand the requirements of such a technology.

15.2.1 *Perception of Marsh Gasses*

While many of the changes that Tidmarsh undergoes are visible to us, some remain outside of our perceptive horizon. For example, marshes produces gasses with unique chemical compositions [176]. I would like to build shoes that make me experience the amount and consistency of marsh-gas in the air, as I walk through Tidmarsh. For simplicity I will limit the example to methane.

The imagined shoe is controlled by a microcomputer, such as a RaspberryPi[1] or a smartphone. This microcomputer is used to find the user's location – using either GPS, or triangulation from sensor nodes or both – and serves as a UI for setting system parameters (the user might want to adjust the type of feedback, or the signal source). The microcomputer also connects to the sensor nodes to collect relevant data on the user's location.

The imagined shoe also has an augmented insole with at least one vibrotactile feedback device controlled by a microcontroller capable of operating at high frequencies and in real time. The microcontroller has access to sensor-data providing information on the movement of the shoe and the amount of pressure exerted on the sole of the shoe. The pressure sensor has a high refresh rate (>1000 Hz) and is sampled with high resolution (>12bit). The IMU is sampled at similarly high speeds. To prevent drift and improve the resolution, multiple IMUs might be used together.

The microcomputer sends high level control-information (amplitude, start, stop, etc.) to the microcontroller. The microcontroller implements these instructions and generates the specific haptic feedback patterns based on the real-time sensor data it has access to. A general overview of the information flow of such a system is presented in Figure 65.

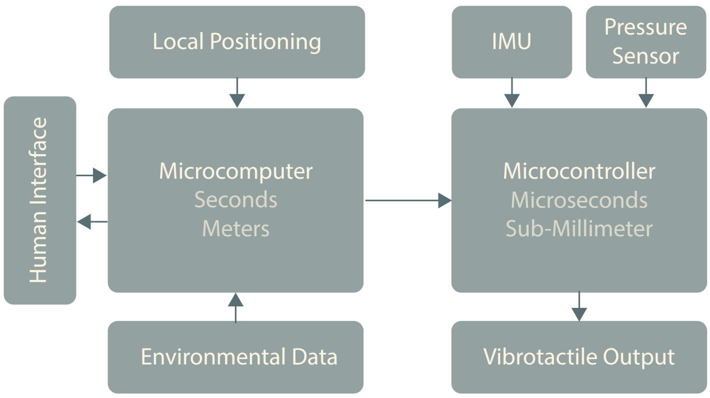

Figure 65: Overview of Information Flow of Haptic System. Note that the left-hand and right-hand components operate on different spatial and temporal scales. The left hand part operates in what I refer to as the *cognitive band* and supports explicit user input. The right hand side operates in the *perceptive band*. The interactions with this part of the system cannot be reduced to discrete input and output – instead it supports perceptive acts which allow experiences to emerge.

1 https://www.raspberrypi.org/

15.2.2 *Cognitive Band*

The cognitive band is where we speak of interactions as consisting of input and output. In my example, I want to design a technology which indicates to the user if methane is present and to what extent. In this context input is *walking to a new location* and output is provided through *vibrotactile feedback relative to the level of methane*. To do so we need to know where the user is and the gas concentration at that location. This is done by the microcomputer and need not happen in real time. The updates could happen several times a second, and the precise location of the user will have a measurement error in the magnitude of meters.

Figure 66 shows aspects of the systems behavior. As the user walks through Tidmarsh the levels of methane vary (progress on the user's path is indicated as a straight line, the x-axis on the left graph, Figure 66). Based on a preset mapping the amount of methane determines the intensity of the haptic feedback (which could be user adjustable, some potential mappings are shown in Figure 66, right). The microcomputer sends this intensity level to the microcontroller, which adjusts the haptic feedback accordingly.

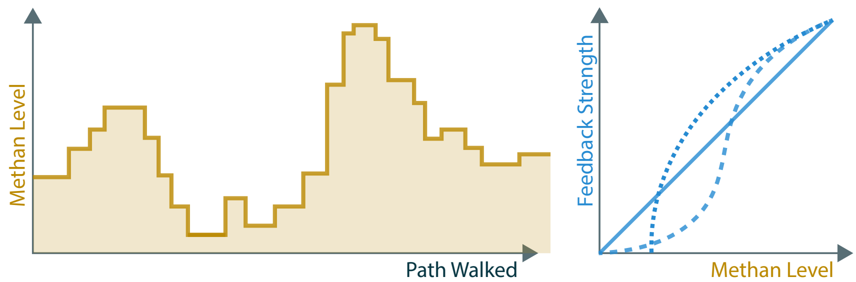

Figure 66: Interaction in the *Cognitive Band* as the user walks through Tidmarsh, the levels of methan vary (left). Based on the methan concentration the amplitude of the vibrotactile feedback is adjusted (right).

15.2.3 *Perceptive Band*

If we would remain in the *cognitive band*, the technology would perform its function without issue. However the vibrotactile feedback would remain purely symbolic. Focussing on the *perceptual band*, we can consider how we want the user to perceive the information.

By closely coupling the vibrotactile feedback to the user's motion, we can support a perception shift and enable the motion of the user and the tactile feedback to merge into a perceptive act. We do this by measuring the movements of the user with high precision. Whenever the measured movement exceeds the threshold set by the granularity, we provide a tactile impulse (see Figure 67). Within the perceptive band, speaking of discrete input or output pairings is not meaningful, instead we can speak of input mappings and describe how they relate to output parameters.

INPUT MAPPINGS Section 13.2 provides an overview of some possible input mappings. For the augmented shoe, I chose isometric pressure sensing, which will enable the experience of compliance to emerge and isotonic motion, which will allow a sensation of resistance to emerge. This means we can change the experienced hardness of the surface the user is walking on and create the experience of having to move ones feet through a material medium.

OUTPUT PARAMETERS Section 13.1 provides an overview of some output parameters we can consider. To generate a vibrotactile signal at all, we must first decide upon how many haptic pulses are provided, given a certain amount of motion – I refer to this as the *granularity* of the signal (see Figure 67). Once we know when the pulses should be created, we must also decide what type of pulse – a square pulse feels very different from a sine. This difference in experience is what I refer to as *timbre*.

One might wish to provide information on other gasses such as hydrogen sulfide or carbon dioxide. This could be done by giving each gas a unique granularity and timbre combination. When multiple gasses are present at the same time the resulting signals might be added, resulting in complex stimulation parameters.

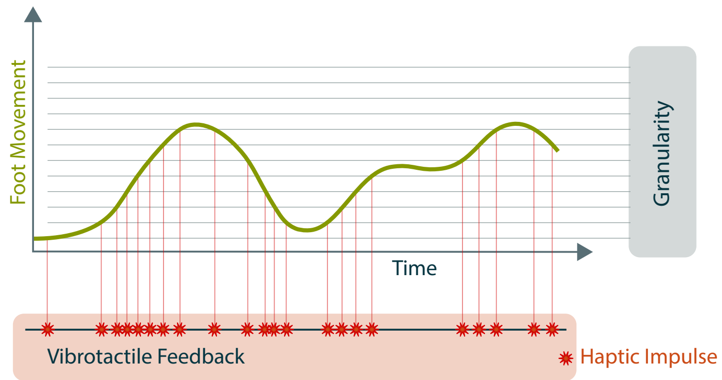

Figure 67: Graphical representation of the isotonic element of proposed system. As the foot moves, it speeds up and slows down based on ones cadence. The change in motion relative to the selected granularity level creates the vibrotactile signal: when the change is large enough to cross a granularity level, a haptic impulse is created. Together, these discrete impulses merge to the motion-coupled vibrotactile feedback through which the material experience is mediatied.

Using the input mappings and output parameters we can design a system where the user experiences themselves as moving through the various gasses as they walk from location to location. The user might literally feel themselves moving through a complex medium which they previously did not know was there. The information is now no longer presented to the user as vibrotactile symbol, but the information is now encoded in new material properties of the world.

15.2.4 *Evaluation*

There are a number of ways in which one might evaluate how this imagined system performs in the cognitive band. For example, we might design an experiment that teaches us if the user is correctly interpreting the vibrotactile signal, or if the user remembers this type of information better than a visual representation. However, such evaluations do not capture the *embodied* dimension of perception, as the above questions are orthogonal to the quality of the mediation. In fact, we are so good at manipulating symbols in the cognitive band that it is quite reasonable that a purely symbolic system performs better than one that provides information using the perceptive band.

When designing for the perceptive band, our goal is in providing the user with a certain quality to the interaction, not a technology which has the highest possible practical utility. We therefore need to evaluate it according to different standards. While more work on this is required, I would currently suggest a multi-step procedure.

Once the considerations of Chapter 14 are organized to a scientific theory, we should begin to understand how the preconditions and temporal factors relate to qualities of the perception shift.

- Check if the preconditions for the perception shift are met (see *Required Conditions* in section 14.4.1)
- Check if sampling rate, latency and jitter are within an acceptable range (see *Boundaries* in section 14.4.1).
- Establish the quality of the perception shift (See *Stability of Perception Shift* in section 14.4.1). This might be done by investigating the following questions:

  Q1 Does the perception shift occur?

  Q2 Is the perception shift persistant?

  Q3 Is the perception shift reverseable at will?

  Q4 Is the perception shift reverseable at all?

  Q5 Is the perception shift multistable?

Assuming that perceptual shifts can have different levels of salience, we might use such an evaluation so assess the extent to which our system truly is capable of embodied mediation of information. Currently this is still speculation, but it is my hope that an empirical theory of embodied interaction might provide us with such a tool.

15.3 CONCLUSION

This Chapter described how the results from Chapter 13 and the considerations from Chapter 14 might be applied in a specific technology.

As suggested in Chapter 14, considering the cognitive and perceptual dimension separately helps us think about the architecture and system requirements of our system. In the cognitive band, we can design interactions considering explicit user input (such as choosing a amplitude mapping) and react accordingly. In the perceptive band, such discrete input and output pairs are no longer relevant. Instead we discuss input mappings and output parameters. Chapter 13 provides some guidance in selecting these.

If our goal is to design for the perceptive band, our evaluation should reflect how well we were able to act upon our design intent. Traditional metrics are therefore not well suited for the evaluation. The theoretical considerations of Chapter 14 provide some initial guidance on factors that might be considered for evaluation.

15.4 ACKNOWLEDGEMENTS

Thanks to Joe Paradiso for hosting me at the MIT Media Lab and Bryan Mayton and Gershon Dublon for introducing me to Tidmarsh and their respective research, which has helped me in reflecting on my own work.

16 CONCLUSION

My interest in pursuing this research was to learn more about how to design haptic feedback, not as a symbol which is interpreted to indicate some information, but to create a direct experience of the information of interest.

I approached this from a practical, applied perspective, by designing Magnetips, a system which can deliver vibrotactile stimuli via the use of on-body or implanted magnets (Chapters 2 – 4), and ReFlex, a bendable smartphone with which I demonstrated how vibrotactile feedback might be created relative to user motion (Chapters 5, 6).

Simultaneously, I approached the problem from an experience-focused perspective. In Chapter 7 I present three basic parameters – amplitude, granularity, and timbre – which can be used to shape vibrotactile feedback. I investigate how variations in the levels of the parameters lead to changes in how the interaction is experienced. The experience of the interaction is formed not only by the feedback parameters, but also by the motion used to generate the vibrations. I demonstrate this in Chapters 10 and 11.

The perception studies teach us about how to design haptic devices, and the extent to which the methods I suggest can be implemented on off-the-shelf devices. For example, because of the importance of timbre, LRAs severely limit the expressivity of haptic devices. On the other hand, the internal IMUs of most mobile devices can be used to capture motion for creating compelling material experiences.

The importance of the experiments is not only in teaching us how to design the next generation of haptic devices, but also teaching us about experience itself. Taking a meta view, the specific outcomes of the experiments are less relevant than the fact that we could collect consistent data in the first place. Participants in the *Haptic Textures* experiment (Chapter 7) consistently reported on variations in material experience even though the actual material of the slider was never altered. In the *Pulse Trains* experiment (Chapter 10), participants provided descriptions of how motion and vibration merge into material experiences, enabling us to reason about the general mechanics behind the phenomenon and pointing towards an empirical approach to embodied interaction (see discussion in Chapter 14). Specifically, I suggest that, when thinking of interaction as occurring in multiple temporal bands, we should consciously consider a perceptive band, occupying the time scale between the cognitive and biological bands. Designing specifically for micro-scale interactions which occur in the perceptive band allows us to present information to the user in a more direct, embodied manner.

Understanding this level of design better on a theoretical level might one day also present us with new tools for evaluating such systems.

I present an imagined device in Chapter 15. This chapter uses the technical knowledge accumulated during this research process and explains how the theoretical considerations presented in the final section of the thesis might help us design systems which provide access to the experience of interest, rather than mediating proxy symbols.

The work presented in this thesis is multi-disciplinary not only in how I ask questions and approach problems, but also in my approach towards evaluation. The thesis includes technical evaluations, user performance studies, quantitative studies of perception, and qualitative in-depth interview studies. Their order in this thesis is not random. The questions that I ask start very narrow and gradually become more open. In *Magnetips*, I ask, *Does it work?* In *ReFlex*, my main question is, *How well can it be used?* In my studies of perception, in *Generating Haptic Textures*, I ask, *How is it experienced?* Finally, in *Pulse Trains*, I ask, *How do people reflect on their experience?*

Wendy Ju mentioned to me that research can be seen as funnels which might either act to open up a space worth exploring, or to hone in on a specific question.

I like to compare this to initially making a tiny hole in a wall, and then gradually breaking more and more away to have a better view on the landscape which is revealed behind it. The landscape here is the relationship between body, motion, and haptic perception. The different levels of breadth provide different utility. The narrow point of departure provides direct utility to designing and implementing related devices. The broad view that results from the very open interview and analysis process helps us to imagine what might be included in a theoretical grounding for future work. Going forward, the next steps should include narrowing the scope again, removing what is not relevant, and testing and refining what remains.

In discussing my research, I have presented material outside of the scope of the specific papers that were published. This has been done in part to show how the included papers connect, but also because I wish to share the material in the hope that it might be useful to others. In-vivo devices for HCI have been little explored, and I hope that the chapter that follows *Magnetips* demonstrates how even a device as simple as an implanted magnet is full of potential. The comments on magnitude estimation in the chapter that reflects on *Haptic Textures* are presented in the hope that they might help others when thinking about methodological issues in their studies of perception. Finally, I hope that the thoughts on perception in general, and the concepts I present in Part III, might be the beginning of a conversation that works towards a stronger theoretical grounding for approaches that use the body in HCI.

GLOSSARY

CHI – Conference on Human Factors in Computing. 5, 102, 173

Confidence Interval – A measure of variability. A 95% confidence interval is designed in such a way that – assuming we repeated the experiment an infinite amount of times – 95% of the time the confidence interval would contain the true mean. Strictly speaking, taking a frequentist perspective, we cannot make any claims about the true mean, aside from stating that the interval is designed such, that we are 95% confident the true mean is contained within. From a Bayesian perspective, we might say that all results within the interval are compatible with the results we found. 23, 104, 106, 109, 111, 112, 115, 119–121, 163, 164

Eccentric Rotating Mass vibration motor (ERM) – A DC motor with an offset (non-symmetric) mass attached to the shaft. As the ERM rotates, a centrifugal force is created, that causes a displacement of the motor. With a high number of revolutions per minute, the motor is constantly being displaced and moved by these asymmetric forces. This is perceived as a vibration. 78, 127, 209

Elastic – In the context of isometric and isotonic input, elastic refers to all input methods which combine both isometric and isotonic properties. 47, 167, 168, 170

Electric Muscle Stimulation (EMS) – The elicitation of muscle contraction using electric impulses. 5, 6

Embodied – The way in which we engage with the world. From a material monist perspective, our consciousness of the world and our engagement with the world are physical activities shaped and performed by our material bodies. I use the term embodied in combination with interaction or perception to highlight that I am explicitly also referring to all the nitty gritty details of these physical activities. This is opposed to more common usage of these terms which are typically based in a cognitivist interpretation, which does not engage with the physicality of existence. I also use the term embodied to refer to a level of understanding the world which does not require symbolic interpretation (see embodied mediation). When I use the term embodied as quoted from somewhere else, it might refer to a plethora of things. An overview of some of these perspectives is provided by van Dijk [184]. 3–6, 178–182, 202

Embodied Mediation – Embodied mediation occurs when a technology mediates one's experience of the world by extending one's access to the world. The technology is interjected between the body and the world in such a way, that one can experience the world through that technology. A typical example is a probe which allows us to reach and feel places we otherwise could not. 3, 147, 182, 195, 202, 207

Experience – I think of experience as the physical state of the body resulting from perception. A less radical view might be that experience is the mental activity emerging from perception. 1–4, 82, 83, 85, 97, 102, 106, 109, 116–118, 125–129, 133, 141, 174, 175, 177, 191, 195

Fitts' Law – A law used in HCI, stating that the time required to rapidly move to a target area is a function of the ratio between the distance to the target and the width of the target. 48, 53, 65, 68, 173

Granularity – I use granularity to describe properties of the haptic signal in the time domain. Specifically, granularity describes how frequently pulses occur based on user action. In translation-based mappings I describe granularity using pulses per centimeter. A sub-dimension of granularity which I do not explore is regularity. Other parameters which might be considered in the time domain include the envolope of haptic signals (attack, decay, sustain, release). 75, 76, 79, 81–83, 88, 90, 91, 93–95, 97, 103, 108, 111, 115, 117, 120, 122, 123, 144, 146, 159, 163–165, 200, 201

Hermeneutic Mediation – Hermeneutic mediation occurs when a technology mediates one's experience of the world by reconfiguring the world. The technology is interjected between the body and the world in such a way, that one is provided with new information which one might interpret. A typical example is a thermometer which presents the temperature to us. 3, 147, 182, 195

Human Computer Interaction (HCI) – A term popularized by Card, Newall, and Moran [28]. In Academia HCI is a discipline that researches the design and use of computing technology. 3–6, 97, 156, 165, 173, 174, 179, 180, 185–187, 189, 190, 208

Inertial Measurement Unit (IMU) – A sensing device which typically consists of an accelerometer, a gyrscope and a magnetometer. Using sensor-fusion, measures from the individual sensor can be combined to determine the orientation of the device. 14, 16, 19, 40, 159, 199, 205

Isometric – Isometric contractions occur when the joint angle and muscle length do not change during contraction. Devices that require isometric contractions (for example, some joysticks or the Thinkpad trackpoint) are typically referred to as isometric input devices. 47, 165–167, 170, 171, 207

Isotonic – Isotonic muscle contractions maintain constant tension in the muscle as the muscle changes length. Devices that require isotonic muscle contractions (such as a mouse) can be referred to as isotonic input devices. 47, 165, 167, 169, 171, 207

Linear Resonate Actuator (LRA) – A vibration motor that produces an oscillating force across a single axis. Unlike a DC ERM motor, a linear resonant actuator relies on an AC voltage to drive a voice coil pressed against a moving mass connected to a spring. When LRAs are driven with a frequency other than their resonant frequency, the performance and efficiency is dramatically reduced. 123, 205

Material Experience – An experience which we associate with physical material. This includes friction, texture, weight, roughness, traction, resistance, compliance etc. I use this term as a 'catch all' phrase for the many experiences which can emerge by varying input mappings and output parameters when coupling vibration to human actions. 1, 2, 4, 63–66, 69, 70, 115, 122, 123, 125, 126, 140–142, 147, 163, 169, 170, 176, 177, 189, 191, 201, 205, 209

NIME – Conference on New Instruments for Musical Expression. 6

Non-Grounded Haptic Feedback – Grounded haptic feedback devices are linked to a fixed point via a kinematic chain. They can provide counter force. Non-grounded haptic feedback devices, on the other hand, are free to move. With some notable exceptions, such as gyroscopic force feedback devices [120], they do not provide counter force. Pachierotti provides an insightfull overview (see [135]). The devices I use are either non-grounded, or do not use the grounding for creating material experiences. 1, 127

Perception – When I speak of perception, I am referring to the actions the body performs from which experiences emerge. It is impossible to feel a texture without relative motion between the body and the texture. Only once the material and the body touch and when relative motion is presence, is the skin vibrated. This vibration then leads to activity of mechanoreceptors. Over time this allows an experience to emerge. When I speak of perception, I refer to this interplay between muscle activity and resulting action potentials emerging from mechanoreceptor which need to occur for us to experience the texture. 1, 3, 4, 77, 137, 164, 173–178, 195, 196

Perception-shift – Under certain conditions, if one is exposed to vibration at a speed relative to the action one is performing, these are perceived not as vibration, but as a material property of the object one is interacting with, or of the objects environment. I refer to this switch from the experience of vibration to the experience of a material property as perception-shift.. 125, 126, 140–142, 147, 154, 169, 188, 189, 191, 196

Pulse – A rapid, transient change in the amplitude of a signal from a baseline value to a higher or lower value, followed by a rapid return to the baseline value. 12, 17, 21, 23, 76, 79, 81, 83, 94, 103, 128–130, 132, 144, 164, 165

Pulse Train – A non-sinusoidal waveform that includes square waves (duty cycle of 50%) and similarly periodic but asymmetrical waves (duty cycles other than 50%). It is a term common to synthesizer programming, and is a typical waveform available on many synthesizers. 76, 83, 126, 128, 130, 132, 140

Quality – I use the term quality for referring to the features of distinct qualia; when reflecting on the difference in experience between blue or red, it is their quality that makes the experiences distinct. 81, 102, 104, 165, 185

Salience – I use salience to describe how strongly a stimulus is experienced. Comparing the color of two blue post-it notes, one forgotten in a corner, and one placed in front of me on my desk, I might perceive them to have the same qualia, but different salience. 102, 104, 165

TEI – Conference on Tangible, Embedded and Embodied Interaction. 5, 6

Timbre – I use timbre as the catch-all term to describe the design parameters of haptic-pulses in the frequency domain. The timbre can be influenced based on how haptic pulses are created: In *Magnetips*, we use pulse-bursts consists of a series of pulses, changing the duration of these pulses changes the frequency of the pulse-burst. Timbre can also be influenced by filtering the signal after it is created. In *Haptic Textures* pulses are passed through a band-pass filter with a narrow Q factor. Changes in timbre are created by changing the central frequency of the filter. 75, 76, 81–83, 85, 86, 88, 90, 91, 93–95, 103, 104, 108, 110, 111, 115, 117, 118, 120, 123, 127, 144, 159, 163–165, 201

[1] Ehud Ahissar and Eldad Assa. "Perception as a closed-loop convergence process." In: *eLife* 5.MAY2016 (2016). ISSN: 2050084X. DOI: 10.7554/eLife.12830 (Cited on pages 185, 188).

[2] Teemu Tuomas Ahmaniemi, Vuokko Lantz, and Juha Marila. "Perception of Dynamic Audiotactile Feedback to Gesture Input." In: *Proceedings of the 10th International Conference on Multimodal Interfaces*. ICMI '08. Chania, Crete, Greece: ACM, 2008, pp. 85–92. ISBN: 978-1-60558-198-9. DOI: 10.1145/1452392.1452410. URL: http://doi.acm.org/10.1145/1452392.1452410 (Cited on page 128).

[3] Teemu Ahmaniemi, Vuokko Lantz, and Juha Marila. "Dynamic Audiotactile Feedback in Gesture Interaction." In: *Proceedings of the 10th International Conference on Human Computer Interaction with Mobile Devices and Services*. MobileHCI '08. Amsterdam, The Netherlands: ACM, 2008, pp. 339–342. ISBN: 978-1-59593-952-4. DOI: 10.1145/1409240.1409283. URL: http://doi.acm.org/10.1145/1409240.1409283 (Cited on page 128).

[4] Motoyuki Akamatsu, Scott Mackenzie, and Thierry Hasbroucq. "A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device." In: *Ergonomics* 38.4 (1995). PMID: 7729406, pp. 816–827. DOI: 10.1080/00140139508925152. eprint: https://doi.org/10.1080/00140139508925152. URL: https://doi.org/10.1080/00140139508925152 (Cited on pages 48, 54, 57).

[5] Michel Amberg, Frédéric Giraud, Betty Semail, Paolo Olivo, Géry Casiez, and Nicolas Roussel. "STIMTAC: A Tactile Input Device with Programmable Friction." In: *Proceedings of the 24th Annual ACM Symposium Adjunct on User Interface Software and Technology* (2011), pp. 7–8. DOI: 10.1145/2046396.2046401 (Cited on page 78).

[6] *Apple's 'force touch' and 'taptic engine' explained.* 2015. URL: https://www.theguardian.com/technology/2015/mar/11/apples-force-touch-taptic-engine-explained-haptic-technology (visited on 09/21/2016) (Cited on page 95).

[7] Daniel Ashbrook, Patrick Baudisch, and Sean White. "Nenya: subtle and eyes-free mobile input with a magnetically-tracked finger ring." In: *Proceedings of the SIGCHI Conference on Hu-

man Factors in Computing Systems. ACM. 2011, pp. 2043–2046 (Cited on page 16).

[8] Matthias Baldauf, Sebastian Zambanini, Peter Fröhlich, and Peter Reichl. "Markerless Visual Fingertip Detection for Natural Mobile Device Interaction." In: *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*. MobileHCI '11. Stockholm, Sweden: ACM, 2011, pp. 539–544. ISBN: 978-1-4503-0541-9. DOI: 10.1145/2037373.2037457. URL: http://doi.acm.org/10.1145/2037373.2037457 (Cited on page 29).

[9] Jared M Batterman and Bruce N Walker. "Displaying Error & Uncertainty in Auditory Graphs." In: (2012), pp. 285–286. DOI: 10.1037/1076-898X.8.4.211 (Cited on pages 97–99, 101).

[10] Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. "TeslaTouch: Electrovibration for Touch Surfaces." In: *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology*. UIST '10. New York, New York, USA: ACM, 2010, pp. 283–292. ISBN: 978-1-4503-0271-5. DOI: 10.1145/1866029.1866074. URL: http://doi.acm.org/10.1145/1866029.1866074 (Cited on pages 46, 48).

[11] Patrick Baudisch and Gerry Chu. "Back-of-device Interaction Allows Creating Very Small Touch Devices." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '09. Boston, MA, USA: ACM, 2009, pp. 1923–1932. ISBN: 978-1-60558-246-7. DOI: 10.1145/1518701.1518995. URL: http://doi.acm.org/10.1145/1518701.1518995 (Cited on page 29).

[12] Sliman J Bensmaia and Mark Hollins. "The vibrations of texture." In: *Somatosensory and Motor Research* 20.1 (2003), pp. 33–43. ISSN: 08990220. DOI: 10.1080/0899022031000083825. URL: http://www.ncbi.nlm.nih.gov/pubmed/12745443http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2074877 (Cited on pages 23, 24, 35, 76, 77, 94, 176, 187).

[13] Sliman Bensmaïa and Mark Hollins. "Pacinian representations of fine surface texture." In: *Perception & Psychophysics* 67.5 (2005), pp. 842–854. ISSN: 1532-5962. DOI: 10.3758/BF03193537. URL: https://doi.org/10.3758/BF03193537 (Cited on pages 76, 77, 94, 125–127, 129).

[14] Sliman Bensmaïa, Mark Hollins, and Jeffrey Yau. "Vibrotactile intensity and frequency information in the pacinian system: a psychophysical model." In: *Perception & psychophysics* 67.5 (2005), pp. 828–841. ISSN: 0031-5117. DOI: 10.3758/BF03193536 (Cited on page 77).

[15] Wouter M. Bergmann Tiest and Astrid M. L. Kappers. “Cues for Haptic Perception of Compliance.” In: *EEE Trans. Haptics* 2.4 (Oct. 2009), pp. 189–199. ISSN: 1939-1412. DOI: 10.1109/TOH.2009.16. URL: http://dx.doi.org/10.1109/TOH.2009.16 (Cited on pages 127, 187).

[16] Joanna Bergstrom-Lehtovirta, Sebastian Boring, and Kasper Hornbæk. “Placing and Recalling Virtual Items on the Skin.” In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI ’17*. New York, New York, USA: ACM Press, 2017, pp. 1497–1507. ISBN: 9781450346559. DOI: 10.1145/3025453.3026030. URL: http://dl.acm.org/citation.cfm?doid=3025453.3026030 (Cited on page 11).

[17] Joanna Bergstrom-Lehtovirta, Kasper Hornbæk, and Sebastian Boring. “It’s a Wrap: Mapping On-Skin Input to Off-Skin Displays.” In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI ’18. Montreal QC, Canada: ACM, 2018, 564:1–564:11. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3174138. URL: http://doi.acm.org/10.1145/3173574.3174138 (Cited on page 5).

[18] Joanna Bergstrom-Lehtovirta, David Coyle, Jarrod Knibbe, and Kasper Hornbæk. “I Really Did That: Sense of Agency with Touchpad, Keyboard, and On-skin Interaction.” In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI ’18. Montreal QC, Canada: ACM, 2018, 378:1–378:8. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173952. URL: http://doi.acm.org/10.1145/3173574.3173952 (Cited on page 5).

[19] Frank Biocca, Jin Kim, and Yung Choi. “Visual Touch in Virtual Environments: An Exploratory Study of Presence, Multimodal Interfaces, and Cross-Modal Sensory Illusions.” In: *Presence: Teleoper. Virtual Environ.* 10.3 (June 2001), pp. 247–265. ISSN: 1054-7460. DOI: 10.1162/105474601300343595. URL: http://dx.doi.org/10.1162/105474601300343595 (Cited on page 48).

[20] Laurens Boer, Ben Cahill, and Anna Vallgårda. “The Hedonic Haptics Player: A Wearable Device to Experience Vibrotactile Compositions.” In: *Proceedings of the 2017 ACM Conference Companion Publication on Designing Interactive Systems*. DIS ’17 Companion. Edinburgh, United Kingdom: ACM, 2017, pp. 297–300. ISBN: 978-1-4503-4991-8. DOI: 10.1145/3064857.3079178. URL: http://doi.acm.org/10.1145/3064857.3079178 (Cited on page 139).

[21] Sebastian Boring, David Ledo, Xiang ’Anthony’ Chen, Nicolai Marquardt, Anthony Tang, and Saul Greenberg. “The Fat Thumb: Using the Thumb’s Contact Size for Single-handed

Mobile Interaction." In: *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services*. MobileHCI '12. San Francisco, California, USA: ACM, 2012, pp. 39–48. ISBN: 978-1-4503-1105-2. DOI: 10.1145/2371574.2371582. URL: http://doi.acm.org/10.1145/2371574.2371582 (Cited on page 11).

[22] Albert S. Bregman. *Auditory scene analysis: the perceptual organization of sound*. MIT Press, 1990, p. 773. ISBN: 9780262022972 (Cited on page 183).

[23] Bret Victor and Bret Victor. *A brief rant on the future of interaction design*. 2011. URL: http://worrydream.com/ABriefRantOnTheFutureOfInteractionDesign/ (visited on 09/21/2016) (Cited on page 95).

[24] Daniel J Brewer, David J Meyer, Michael A Peshkin, and J Edward Colgate. "Viscous Textures: Velocity Dependence in Fingertip-Surface Scanning Interaction." In: *Haptics Symposium, 2016 IEEE* (2016), pp. 0–5. ISSN: 23247355. DOI: 10.1109/HAPTICS.2016.7463188 (Cited on page 94).

[25] Lorna M Brown and Joseph Jofish Kaye. "Eggs-ploring the influence of material properties on haptic experience." In: *HAID'07* (2007) (Cited on page 128).

[26] Jesse Burstyn, Juan Pablo Carrascal, and Roel Vertegaal. "Fitts' Law and the Effects of Input Mapping and Stiffness on Flexible Display Interactions." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: ACM, 2016, pp. 3649–3658. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858383. URL: http://doi.acm.org/10.1145/2858036.2858383 (Cited on page 54).

[27] Alex Butler, Shahram Izadi, and Steve Hodges. "SideSight: Multi-"touch" Interaction Around Small Devices." In: *Proceedings of the 21st annual ACM symposium on User interface software and technology - UIST '08*. New York, New York, USA: ACM Press, 2008, p. 201. ISBN: 9781595939753. DOI: 10.1145/1449715.1449746. URL: http://dl.acm.org/citation.cfm?id=1449715.1449746 (Cited on page 14).

[28] Stuart K Card, Thomas P Moran, and Allen Newell. *The psychology of human-computer interaction. 1983*. 1983 (Cited on pages 186, 208).

[29] Taylor Carman. "The Body in Husserl and Merleau-Ponty." In: *Philosophical Topics* 27.2 (1999), pp. 205–226 (Cited on page 193).

[30] Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. "UltraHaptics: Multi-point Midair Haptic Feedback for Touch Surfaces." In: *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*. UIST '13. St. Andrews, Scotland, United Kingdom: ACM, 2013, pp. 505–514. ISBN: 978-1-4503-2268-3. DOI: 10.1145/2501988.2502018. URL: http://doi.acm.org/10.1145/2501988.2502018 (Cited on pages 15, 127, 128).

[31] Géry Casiez and Daniel Vogel. "The Effect of Spring Stiffness and Control Gain with an Elastic Rate Control Pointing Device." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '08. Florence, Italy: ACM, 2008, pp. 1709–1718. ISBN: 978-1-60558-011-1. DOI: 10.1145/1357054.1357321. URL: http://doi.acm.org/10.1145/1357054.1357321 (Cited on page 47).

[32] Ke-Yu Chen, Shwetak N Patel, and Sean Keller. "Finexus: Tracking precise motions of multiple fingertips using magnetic sensing." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM. 2016, pp. 1504–1514 (Cited on pages 15–17).

[33] Ke-Yu Chen, Kent Lyons, Sean White, and Shwetak Patel. "uTrack: 3D input using two magnetic sensors." In: *Proceedings of the 26th annual ACM symposium on User interface software and technology*. ACM. 2013, pp. 237–244 (Cited on page 16).

[34] Nicholas Chen, Francois Guimbretiere, Morgan Dixon, Cassandra Lewis, and Maneesh Agrawala. "Navigation Techniques for Dual-display e-Book Readers." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '08. Florence, Italy: ACM, 2008, pp. 1779–1788. ISBN: 978-1-60558-011-1. DOI: 10.1145/1357054.1357331. URL: http://doi.acm.org/10.1145/1357054.1357331 (Cited on page 60).

[35] Juliana Cherston, Paul Strohmeier, and Joseph A Paradiso. "Grappler: Array of Bistable Elements For Pinching Net-Like Infrastructure to Low Gravity Bodies." In: *AIAA Scitech 2019 Forum*. 2019, p. 0871 (Cited on page 7).

[36] Youngjun Cho, Andrea Bianchi, Nicolai Marquardt, and Nadia Bianchi-Berthouze. "RealPen: Providing Realism in Handwriting Tasks on Touch Surfaces Using Auditory-Tactile Feedback." In: *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. UIST '16. Tokyo, Japan: ACM, 2016, pp. 195–205. ISBN: 978-1-4503-4189-9. DOI: 10.1145/2984511.2984550. URL: http://doi.acm.org/10.1145/2984511.2984550 (Cited on page 2).

[37] Luke Cramphorn, Benjamin Ward-Cherrier, and Nathan F. Lepora. "A biomimetic fingerprint improves spatial tactile perception." In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 9793. Springer, Cham, 2016, pp. 418–423. ISBN: 9783319424163. DOI: 10.1007/978-3-319-42417-0_40. URL: http://link.springer.com/10.1007/978-3-319-42417-0{_}40 (Cited on page 176).

[38] Heather Culbertson and Katherine J. Kuchenbecker. "Importance of Matching Physical Friction, Hardness, and Texture in Creating Realistic Haptic Virtual Surfaces." In: *EEE Trans. Haptics* 10.1 (Jan. 2017), pp. 63–74. ISSN: 1939-1412. DOI: 10.1109/TOH.2016.2598751. URL: https://doi.org/10.1109/TOH.2016.2598751 (Cited on page 127).

[39] Heather Culbertson and Katherine J Kuchenbecker. "Ungrounded haptic augmented reality system for displaying roughness and friction." In: *IEEE/ASME Transactions on Mechatronics* 22.4 (2017), pp. 1839–1849 (Cited on pages 127, 128, 175).

[40] Heather Culbertson, Juliette Unwin, Benjamin E Goodman, and Katherine J Kuchenbecker. "Generating haptic texture models from unconstrained tool-surface interactions." In: *2013 World Haptics Conference (WHC)*. IEEE. 2013, pp. 295–300 (Cited on pages 76, 79, 126–129, 132).

[41] Heather Culbertson, Julie M. Walker, Michael Raitor, and Allison M. Okamura. "WAVES: A Wearable Asymmetric Vibration Excitation System for Presenting Three-Dimensional Translation and Rotation Cues." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: ACM, 2017, pp. 4972–4982. ISBN: 978-1-4503-4655-9. DOI: 10.1145/3025453.3025741. URL: http://doi.acm.org/10.1145/3025453.3025741 (Cited on pages 127, 128).

[42] Artem Dementyev, Javier Hernandez, Inrak Choi, Sean Follmer, and Joseph Paradiso. "Epidermal Robots: Wearable Sensors That Climb on the Skin." In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2.3 (Sept. 2018), 102:1–102:22. ISSN: 2474-9567. DOI: 10.1145/3264912. URL: http://doi.acm.org/10.1145/3264912 (Cited on page 41).

[43] Jelle van Dijk and Caroline Hummels. "Designing for Embodied Being-in-the-World: Two Cases, Seven Principles and One Framework." In: *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '17. Yokohama, Japan: ACM, 2017, pp. 47–56. ISBN: 978-1-4503-4676-4. DOI: 10.1145/3024969.3025007. URL: http://doi.

acm.org/10.1145/3024969.3025007 (Cited on pages 3, 180, 195).

[44] Jelle van Dijk, Remko van der Lugt, and Caroline Hummels. "Beyond Distributed Representation: Embodied Cognition Design Supporting Socio-sensorimotor Couplings." In: *Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction*. TEI '14. Munich, Germany: ACM, 2014, pp. 181–188. ISBN: 978-1-4503-2635-3. DOI: 10.1145/2540930.2540934. URL: http://doi.acm.org/10.1145/2540930.2540934 (Cited on pages 179, 180, 195).

[45] Maurin Donneaud, Cedric Honnet, and Paul Strohmeier. "Designing a multi-touch etextile for music performances." In: *NIME*. 2017, pp. 7–12 (Cited on page 6).

[46] Paul Dourish. *Where the action is: the foundations of embodied interaction*. Massechusetts: The MIT Press, 2001. Chap. 1. ISBN: 0262541785. DOI: 10.1162/leon.2003.36.5.412 (Cited on pages 179, 180, 195).

[47] Gershon Dublon. "Sensor(y) Landscapes : technologies for new perceptual sensibilities." In: (2018). URL: https://dspace.mit.edu/handle/1721.1/119070 (Cited on page 198).

[48] Jamie Ferguson and Stephen A Brewster. "Investigating Perceptual Congruence Between Data and Display Dimensions in Sonification." In: (2018), pp. 1–9 (Cited on pages 98, 99, 101, 111).

[49] Jamie Ferguson, John Williamson, and Stephen Brewster. "Evaluating Mapping Designs for Conveying Data through Tactons." In: (2018) (Cited on pages 97–99, 101, 102, 105).

[50] Moffett Field, Moffett Field, and Stephen R Ellis. "Perception of Image Motion During Head Movement." In: 1.212 (2006), pp. 45–50 (Cited on pages 98, 99).

[51] Eelke Folmer and Tony Morelli. "Spatial Gestures Using a Tactile-proprioceptive Display." In: *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*. TEI '12. Kingston, Ontario, Canada: ACM, 2012, pp. 139–142. ISBN: 978-1-4503-1174-8. DOI: 10.1145/2148131.2148161. URL: http://doi.acm.org/10.1145/2148131.2148161 (Cited on page 144).

[52] Clifton Forlines and Ravin Balakrishnan. "Evaluating Tactile Feedback and Direct vs. Indirect Stylus Input in Pointing and Crossing Selection Tasks." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '08. Florence, Italy: ACM, 2008, pp. 1563–1572. ISBN: 978-1-60558-011-1. DOI: 10.1145/1357054.1357299. URL: http://doi.acm.org/10.1145/1357054.1357299 (Cited on pages 48, 54, 57).

[53] Euan Freeman, Stephen Brewster, and Vuokko Lantz. “Tactile Feedback for Above-Device Gesture Interfaces: Adding Touch to Touchless Interactions.” In: *Proceedings of the 16th International Conference on Multimodal Interaction*. ICMI ’14. Istanbul, Turkey: ACM, 2014, pp. 419–426. ISBN: 978-1-4503-2885-2. DOI: 10.1145/2663204.2663280. URL: http://doi.acm.org/10.1145/2663204.2663280 (Cited on pages 11, 14).

[54] Rachel Freire, Cedric Honnet, and Paul Strohmeier. “Second Skin: An Exploration of eTextile Stretch Circuits on the Body.” In: *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*. ACM. 2017, pp. 653–658 (Cited on page 6).

[55] Rachel Freire, Paul Strohmeier, Cedric Honnet, Jarrod Knibbe, and Sophia Brueckner. “Designing eTextiles for the Body: Shape, Volume and Motion.” In: *ACM International Conference on Tangible, Embedded and Embodied Interaction (TEI)*. 2018 (Cited on pages 5, 6).

[56] A. S. French and P. H. Torkkeli. “Mechanoreceptors.” In: *Encyclopedia of Neuroscience* (2010), pp. 689–695. DOI: 10.1016/B978-008045046-9.01921-5 (Cited on page 36).

[57] John P Frisby and Jeremy L Clatworthy. “Illusory Contours: Curious Cases of Simultaneous Brightness Contrast?” In: *Perception* 4.3 (1975), pp. 349–357. DOI: 10.1068/p040349. eprint: https://doi.org/10.1068/p040349. URL: https://doi.org/10.1068/p040349 (Cited on page 147).

[58] Robert H Gault. “Touch as a substitute for hearing in the interpretation and control of speech.” In: *Archives of Otolaryngology* 3.2 (1926), pp. 121–135 (Cited on pages 78, 128).

[59] Stanley A. Gelfand. *Essentials of audiology*. Thieme, 2001, p. 585. ISBN: 1588900177. URL: https://books.google.de/books/about/Essentials_of_Audiology.html?id=UqXS8CLJML4C&redir_esc=y (Cited on page 102).

[60] George A Gescheider. “Psychophysical Schaling.” In: *Ann. Rev. Psychol* 39 (1988), pp. 169–200 (Cited on pages 23, 82, 97–99, 102, 104).

[61] George A. Gescheider. *Psychophysics : the fundamentals*. L. Erlbaum Associates, 1997, p. 435. ISBN: 080582281X (Cited on pages 85, 87, 99, 101, 102, 106).

[62] Gijs Huisman. *Many hands make light work: reflections on Reshaping Touch Communication at CHI’18*. URL: https://in-touch-digital.com/2018/06/28/many-hands-make-light-work-reflections-on-reshaping-touch-communication-at-chi18-guest-post-by-gijs-huisman/ (visited on 03/14/2019) (Cited on page 160).

[63] Iain D. Gilchrist, Valerie Brown, and John M. Findlay. "Saccades without eye movements." In: *Nature* 390.6656 (1997), pp. 130–131. ISSN: 0028-0836. DOI: 10.1038/36478. URL: http://www.nature.com/articles/36478 (Cited on page 193).

[64] Marcello Giordano, Stephen Sinclair, and Marcelo M. Wanderley. *Bowing a vibration-enhanced force feedback device*. University of Michigan, 2012 (Cited on page 78).

[65] Audrey Girouard, Aneesh Tarun, and Roel Vertegaal. "DisplayStacks: Interaction Techniques for Stacks of Flexible Thin-film Displays." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: ACM, 2012, pp. 2431–2440. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2208406. URL: http://doi.acm.org/10.1145/2207676.2208406 (Cited on page 60).

[66] Donald Derek Haddad, Gershon Dublon, Brian D. Mayton, Spencer Russell, Xiao Xiao, Ken Perlin, and Joseph A. Paradiso. "Resynthesizing reality: driving vivid virtual environments from sensor networks." In: *Special Interest Group on Computer Graphics and Interactive Techniques Conference, SIGGRAPH 2017, Los Angeles, California, USA, July 30 - August 3, 2017, Talks*. 2017, 51:1–51:2. DOI: 10.1145/3084363.3085027. URL: https://doi.org/10.1145/3084363.3085027 (Cited on page 198).

[67] Elisabet Hagert. "Proprioception of the Wrist Joint: A Review of Current Concepts and Possible Implications on the Rehabilitation of the Wrist." In: *Journal of Hand Therapy* 23.1 (2010), pp. 2–17. ISSN: 0894-1130. DOI: 10.1016/J.JHT.2009.09.008. URL: https://www.sciencedirect.com/science/article/pii/S0894113009001288 (Cited on page 36).

[68] Sung H. Han, Maengkee Song, and Jiyoung Kwahk. "A systematic method for analyzing magnitude estimation data." In: *International Journal of Industrial Ergonomics* 23.5 (1999), pp. 513–524. ISSN: 01698141. DOI: 10.1016/S0169-8141(98)00017-1 (Cited on pages 87, 101, 106).

[69] Xinying Han, Hiroaki Seki, and M Hikizu. "Wearable handwriting input device using magnetic field." In: *SICE, 2007 Annual Conference*. IEEE. 2007, pp. 365–368 (Cited on page 16).

[70] Chris Harrison and Scott E. Hudson. "Abracadabra." In: *Proceedings of the 22nd annual ACM symposium on User interface software and technology - UIST '09* (2009), p. 121. DOI: 10.1145/1622176.1622199. URL: http://portal.acm.org/citation.cfm?doid=1622176.1622199 (Cited on pages 11, 16, 28).

[71] Chris Harrison, Zhiquan Yeo, and Scott E. Hudson. "Faster Progress Bars: Manipulating Perceived Duration with Visual Augmentations." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '10. Atlanta, Georgia, USA: ACM, 2010, pp. 1545–1548. ISBN: 978-1-60558-929-9. DOI: 10.1145/1753326.1753556. URL: http://doi.acm.org/10.1145/1753326.1753556 (Cited on page 55).

[72] Stefan Helmreich. "Gravity's Reverb: Listening to Space-Time, or Articulating the Sounds of Gravitational-Wave Detection." In: *Cultural Anthropology* 31.4 (2016), pp. 464–492. ISSN: 08867356. DOI: 10.14506/ca31.4.02 (Cited on page 1).

[73] Ken Hinckley, Seongkook Heo, Michel Pahud, Christian Holz, Hrvoje Benko, Abigail Sellen, Richard Banks, Kenton O Hara, Gavin Smyth, and Bill Buxton. "Pre-Touch Sensing for Mobile Interaction." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016), pp. 2869–2881. ISSN: 2334-2536. DOI: 10.1145/2858036.2858095. arXiv: 1605.09782. URL: http://dl.acm.org/citation.cfm?doid=2858036.2858095 (Cited on page 14).

[74] Masaharu Hirose, Karin Iwazaki, Kozue Nojiri, Minato Takeda, Yuta Sugiura, and Masahiko Inami. "Gravitamine Spice: A System That Changes the Perception of Eating Through Virtual Weight Sensation." In: *Proceedings of the 6th Augmented Human International Conference*. AH '15. Singapore, Singapore: ACM, 2015, pp. 33–40. ISBN: 978-1-4503-3349-8. DOI: 10.1145/2735711.2735795. URL: http://doi.acm.org/10.1145/2735711.2735795 (Cited on page 48).

[75] Eve Hoggan, Yi-Ta Hsieh, Kalle Myllymaa, Vuokko Lantz, Johan Kildal, Julian Eiler, and Giulio Jacucci. "An Exploration of Mobile Shape-Changing Textures." In: *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '17. Yokohama, Japan: ACM, 2017, pp. 275–282. ISBN: 978-1-4503-4676-4. DOI: 10.1145/3024969.3024983. URL: http://doi.acm.org/10.1145/3024969.3024983 (Cited on page 127).

[76] David Holman and Roel Vertegaal. "Organic User Interfaces: Designing Computers in Any Way, Shape, or Form." In: *Commun. ACM* 51.6 (June 2008), pp. 48–55. ISSN: 0001-0782. DOI: 10.1145/1349026.1349037. URL: http://doi.acm.org/10.1145/1349026.1349037 (Cited on page 46).

[77] Kasper Hornbæk, Aske Mottelson, Jarrod Knibbe, and Dan Vogel. "What Do We Mean by 'Interaction'? An Analysis of 35 Years of CHI." In: *ACM Transactions on Computer-Human Interaction*. TOCHI (2019) (Cited on pages 173, 179).

[78] Caroline Hummels and Jelle van Dijk. "Seven Principles to Design for Embodied Sensemaking." In: *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '15. Stanford, California, USA: ACM, 2015, pp. 21–28. ISBN: 978-1-4503-3305-4. DOI: 10.1145/2677199.2680577. URL: http://doi.acm.org/10.1145/2677199.2680577 (Cited on pages 180, 195).

[79] Ali Israr and Ivan Poupyrev. "Control space of apparent haptic motion." In: *2011 IEEE World Haptics Conference* (2011), pp. 457–462. DOI: 10.1109/WHC.2011.5945529. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5945529 (Cited on page 78).

[80] Ali Israr and Ivan Poupyrev. "Tactile Brush: Drawing on Skin with a Tactile Grid Display." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pp. 2019–2028. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1979235. URL: http://doi.acm.org/10.1145/1978942.1979235 (Cited on page 128).

[81] I. Darian-Smith J.W. Morley, A.W. Goodwin. "Tactile Discrimination of Gratings." In: *Experimental Brain Research* 9 (1981), pp. 146–158 (Cited on page 94).

[82] Yvonne Jansen, Thorsten Karrer, and Jan Borchers. "MudPad: Tactile Feedback and Haptic Texture Overlay for Touch Surfaces." In: *ACM International Conference on Interactive Tabletops and Surfaces*. ITS '10. Saarbrücken, Germany: ACM, 2010, pp. 11–14. ISBN: 978-1-4503-0399-6. DOI: 10.1145/1936652.1936655. URL: http://doi.acm.org/10.1145/1936652.1936655 (Cited on page 127).

[83] Programmer Joe. *SteamVR Tracking HDK*. 2016. URL: http://steamcommunity.com/app/507090/discussions/0/360671247404603033/ (visited on 09/21/2016) (Cited on page 94).

[84] Roland S Johansson and J Randall Flanagan. "Coding and use of tactile signals from the fingertips in object manipulation tasks." In: *Nature Reviews Neuroscience* 10.5 (2009), p. 345 (Cited on pages 24, 35, 126, 127, 187, 189).

[85] Lynette A. Jones and Hong Z. Tan. "Application of psychophysical techniques to haptic research." In: *IEEE Transactions on Haptics* 6.3 (2013), pp. 268–284. ISSN: 19391412. DOI: 10.1109/TOH.2012.74 (Cited on pages 82, 87).

[86] Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. "The reacTable: Exploring the Synergy Between Live Music Performance and Tabletop Tangible Interfaces." In: *Proceedings of the 1st International Conference on*

Tangible and Embedded Interaction. TEI '07. Baton Rouge, Louisiana: ACM, 2007, pp. 139–146. ISBN: 978-1-59593-619-6. DOI: 10.1145/1226969.1226998. URL: http://doi.acm.org/10.1145/1226969.1226998 (Cited on page 187).

[87] Azusa Kadomura and Itiro Siio. "MagNail: user interaction with smart device through magnet attached to fingernail." In: *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*. ACM. 2015, pp. 309–312 (Cited on page 32).

[88] David Katz. "The World of Touch." In: *Nature* 170.4333 (1952). Ed. by Lester E. Krueger, pp. 815–816. ISSN: 0028-0836. DOI: 10.1038/170815a0. URL: https://www.nature.com/articles/170815a0 (Cited on pages 77, 125).

[89] Hamed Ketabdar, Mehran Roshandel, and Kamer Ali Yüksel. "Towards using embedded magnetic field sensor for around mobile device 3D interaction." In: *Proceedings of the 12th international conference on Human computer interaction with mobile devices and services*. ACM. 2010, pp. 153–156 (Cited on page 15).

[90] Johan Kildal. "3D-press: Haptic Illusion of Compliance when Pressing on a Rigid Surface." In: *International Conference on Multimodal Interfaces and the Workshop on Machine Learning for Multimodal Interaction*. ICMI-MLMI '10. Beijing, China: ACM, 2010, 21:1–21:8. ISBN: 978-1-4503-0414-6. DOI: 10.1145/1891903.1891931. URL: http://doi.acm.org/10.1145/1891903.1891931 (Cited on pages 2, 46, 52, 76, 79, 122, 127, 128, 130, 164, 166, 175, 190).

[91] Johan Kildal, Susanna Paasovaara, and Viljakaisa Aaltonen. "Kinetic Device: Designing Interactions with a Deformable Mobile Interface." In: *CHI '12 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '12. Austin, Texas, USA: ACM, 2012, pp. 1871–1876. ISBN: 978-1-4503-1016-1. DOI: 10.1145/2212776.2223721. URL: http://doi.acm.org/10.1145/2212776.2223721 (Cited on page 47).

[92] Johan Kildal and Graham Wilson. "Feeling It: The Roles of Stiffness, Deformation Range and Feedback in the Control of Deformable Ui." In: *Proceedings of the 14th ACM International Conference on Multimodal Interaction*. ICMI '12. Santa Monica, California, USA: ACM, 2012, pp. 393–400. ISBN: 978-1-4503-1467-1. DOI: 10.1145/2388676.2388766. URL: http://doi.acm.org/10.1145/2388676.2388766 (Cited on pages 47, 48).

[93] Ryo Kitada, Toshihiro Hashimoto, Takanori Kochiyama, Tomonori Kito, Tomohisa Okada, Michikazu Matsumura, Susan J. Lederman, and Norihiro Sadato. "Tactile estimation of the roughness of gratings yields a graded response in the human brain: An fMRI study." In: *NeuroImage* 25.1 (2005), pp. 90–100. ISSN: 10538119. DOI: 10.1016/j.neuroimage.2004.11.026 (Cited on page 94).

[94] Roberta L. Klatzky and Susan J. Lederman. "The Perceived Roughness of Resistive Virtual Textures: I. Rendering by a Force-feedback Mouse." In: *ACM Trans. Appl. Percept.* 3.1 (Jan. 2006), pp. 1–14. ISSN: 1544-3558. DOI: 10.1145/1119766.1119767. URL: http://doi.acm.org/10.1145/1119766.1119767 (Cited on page 48).

[95] Roberta L. Klatzky, Dianne Pawluk, and Angelika Peer. "Haptic perception of material properties and implications for applications." In: *Proceedings of the IEEE* 101.9 (2013), pp. 2081–2092. ISSN: 00189219. DOI: 10.1109/JPROC.2013.2248691 (Cited on page 77).

[96] Scott R. Klemmer, Björn Hartmann, and Leila Takayama. "How Bodies Matter: Five Themes for Interaction Design." In: *Proceedings of the 6th Conference on Designing Interactive Systems.* DIS '06. University Park, PA, USA: ACM, 2006, pp. 140–149. ISBN: 1-59593-367-0. DOI: 10.1145/1142405.1142429. URL: http://doi.acm.org/10.1145/1142405.1142429 (Cited on pages 3, 180, 195).

[97] Jarrod Knibbe, Diego Martinez Plasencia, Christopher Bainbridge, Chee-Kin Chan, Jiawei Wu, Thomas Cable, Hassan Munir, and David Coyle. "Extending interaction for smart watches." In: *Proceedings of the extended abstracts of the 32nd annual ACM conference on Human factors in computing systems - CHI EA '14.* New York, New York, USA: ACM Press, 2014, pp. 1891–1896. ISBN: 9781450324748. DOI: 10.1145/2559206.2581315. URL: https://dl.acm.org/citation.cfm?doid=2559206.2581315 (Cited on pages 11, 14).

[98] Jarrod Knibbe, Paul Strohmeier, Sebastian Boring, and Kasper Hornbæk. "Automatic calibration of high density electric muscle stimulation." In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1.3 (2017), p. 68 (Cited on pages 5, 6).

[99] Laurens R. Krol, Dzmitry Aliakseyeu, and Sriram Subramanian. "Haptic Feedback in Remote Pointing." In: *CHI '09 Extended Abstracts on Human Factors in Computing Systems.* CHI EA '09. Boston, MA, USA: ACM, 2009, pp. 3763–3768. ISBN: 978-1-60558-247-4. DOI: 10.1145/1520340.1520568. URL: http://doi.acm.org/10.1145/1520340.1520568 (Cited on page 128).

[100] Byron Lahey, Audrey Girouard, Winslow Burleson, and Roel Vertegaal. "PaperPhone: Understanding the Use of Bend Gestures in Mobile Devices with Flexible Electronic Paper Displays." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pp. 1303–1312. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1979136. URL: http://doi.acm.org/10.1145/1978942.1979136 (Cited on page 47).

[101] Mathieu Le Goc, Stuart Taylor, Shahram Izadi, and Cem Keskin. "A low-cost transparent electric field sensor for 3d interaction on mobile devices." In: *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM. 2014, pp. 3167–3170 (Cited on page 14).

[102] Susan J. Lederman and Lynette A. Jones. "Tactile and Haptic Illusions." In: *EEE Trans. Haptics* 4.4 (Oct. 2011), pp. 273–294. ISSN: 1939-1412. DOI: 10.1109/TOH.2011.2. URL: http://dx.doi.org/10.1109/TOH.2011.2 (Cited on page 48).

[103] Susan Lederman. "Tactile roughness of grooved surfaces: The touching process and effects of macro and microsurface structure." In: *Perceptual Psychology* 15 (1989), pp. 45–57 (Cited on page 77).

[104] Vincent Levesque, Louise Oram, Karon MacLean, Andy Cockburn, Nicholas D. Marchuk, Dan Johnson, J. Edward Colgate, and Michael A. Peshkin. "Enhancing Physicality in Touch Interaction with Programmable Friction." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pp. 2481–2490. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1979306. URL: http://doi.acm.org/10.1145/1978942.1979306 (Cited on page 127).

[105] Vincent Lévesque, Louise Oram, Karon MacLean, J. Edward Colgate, and Michael a. Peshkin. "Restoring physicality to touch interaction with programmable friction." In: *Digest of Technical Papers - IEEE International Conference on Consumer Electronics* (2011), pp. 61–62. ISSN: 0747668X. DOI: 10.1109/ICCE.2011.5722773 (Cited on page 78).

[106] Yi-Chi Liao, Sunjun Kim, and Antti Oulasvirta. "One Button to Rule Them All: Rendering Arbitrary Force-Displacement Curves." In: *The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings*. UIST '18 Adjunct. Berlin, Germany: ACM, 2018, pp. 111–113. ISBN: 978-1-4503-5949-8. DOI: 10.1145/3266037.3266118. URL: http://doi.acm.org/10.1145/3266037.3266118 (Cited on pages 188, 197).

[107] G E Loeb and J A Fishel. "The role of fingerprints in vibrotactile discrimination." In: *White paper for DoD Physics of Biology* (2009) (Cited on pages 75, 77, 125, 126, 176).

[108] I. Scott. MacKenzie. *Human-computer interaction.* Elsevier Science, 2012, p. 370. ISBN: 0124058655 (Cited on page 186).

[109] Eddy Maddalena and Stefano Mizzaro. *On Crowdsourcing Relevance Magnitudes for Information.* Vol. 35. 3. 2017. ISBN: 1126900001 (Cited on pages 98, 99).

[110] Sebastian OH Madgwick, Andrew JL Harrison, and Ravi Vaidyanathan. "Estimation of IMU and MARG orientation using a gradient descent algorithm." In: *Rehabilitation Robotics (ICORR), 2011 IEEE International Conference on.* IEEE. 2011, pp. 1–7 (Cited on page 18).

[111] Teemu Mäki-patola and Perttu Hämäläinen. "Latency Tolerance for Gesture Controlled Continuous Sound Instrument without Tactile Feedback." In: *Proceedings of the International Computer Music Conference* 1998 (2004) (Cited on page 190).

[112] Elena Márquez Segura, Laia Turmo Vidal, Asreen Rostami, and Annika Waern. "Embodied Sketching." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.* CHI '16. San Jose, California, USA: ACM, 2016, pp. 6014–6027. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858486. URL: http://doi.acm.org/10.1145/2858036.2858486 (Cited on page 180).

[113] Catherine C. Marshall and Sara Bly. "Turning the Page on Navigation." In: *Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries.* JCDL '05. Denver, CO, USA: ACM, 2005, pp. 225–234. ISBN: 1-58113-876-8. DOI: 10.1145/1065385.1065438. URL: http://doi.acm.org/10.1145/1065385.1065438 (Cited on page 60).

[114] Brian D. Mayton, Gershon Dublon, Spencer Russell, Evan F. Lynch, Donald Derek Haddad, Vasant Ramasubramanian, Clement Duhart, Glorianna Davenport, and Joseph A. Paradiso. "The Networked Sensory Landscape: Capturing and Experiencing Ecological Change Across Scales." In: *Presence* 26.2 (2017). DOI: 10.1162/PRES_a_00292. URL: https://doi.org/10.1162/PRES_a_00292 (Cited on page 198).

[115] Jess McIntosh, Paul Strohmeier, Jarrod Knibbe, Sebastian Boring, and Kasper Hornbæk. "Magnetips: Combining Fingertip Track-ing and Haptic Feedback for Around-Device Interaction." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems.* CHI '19 (2019). DOI: 10.1145/3290605.3300638. URL: https://doi.org/10.1145/3290605.3300638 (Cited on pages 36, 39, 41, 42, 146).

[116] Mick Mcgee and D Ph. "Master Usability Scaling : Magnitude Estimation and Master Scaling Applied to Usability Measurement." In: 6.1 (2004), pp. 335–342 (Cited on pages 97–99, 101).

[117] Maurice Merleau-Ponty. *Phenomenology of Perception*. Trans. by Donald A Landes. Routledge, 2012, p. 606. ISBN: 0203994612. DOI: 10 . 4324 / 9780203981139 (Cited on page 173).

[118] Aske Mottelson, Jarrod Knibbe, and Kasper Hornbæk. "Veritaps: Truth Estimation from Mobile Interaction." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: ACM, 2018, 561:1–561:12. ISBN: 978-1-4503-5620-6. DOI: 10 . 1145 / 3173574 . 3174135. URL: http://doi.acm.org/10.1145/3173574.3174135 (Cited on page 5).

[119] Florian "Floyd" Mueller et al. "Body-centric Computing: Results from a Weeklong Dagstuhl Seminar in a German Castle." In: *Interactions* 25.4 (June 2018), pp. 34–39. ISSN: 1072-5520. DOI: 10.1145/3215854. URL: http://doi.acm.org/10.1145/3215854 (Cited on page 180).

[120] Martin Murer, Bernhard Maurer, Hermann Huber, Ilhan Aslan, and Manfred Tscheligi. "TorqueScreen: Actuated Flywheels for Ungrounded Kinaesthetic Feedback in Handheld Devices." In: *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '15. Stanford, California, USA: ACM, 2015, pp. 161–164. ISBN: 978-1-4503-3305-4. DOI: 10.1145/2677199.2680579. URL: http://doi.acm.org/10.1145/2677199.2680579 (Cited on pages 127, 209).

[121] Ken Nakagaki, Luke Vink, Jared Counts, Daniel Windham, Daniel Leithinger, Sean Follmer, and Hiroshi Ishii. "Materiable: Rendering Dynamic Material Properties in Response to Direct Physical Touch with Shape Changing Interfaces." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2016, pp. 2764–2772. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858104. URL: http://doi.acm.org/10.1145/2858036.2858104 (Cited on page 95).

[122] Yusuke Nakagawa, Akiya Kamimura, and Yoichiro Kawaguchi. "MimicTile: A Variable Stiffness Deformable User Interface for Mobile Devices." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: ACM, 2012, pp. 745–748. ISBN: 978-1-4503-1015-4. DOI: 10.1145/2207676.2207782. URL: http://doi.acm.org/10.1145/2207676.2207782 (Cited on page 47).

[123] Takaaki Nara, Takasaki Masaya, Tachi Susumu, and Higuchi Toshiro. "An Application of SAW to a Tactile Display in Virtual Reality." In: *IEEE Ultrasonics Symposium* (2000), pp. 1–4 (Cited on pages 78, 79, 122).

[124] Allen Newell. *Unified theories of cognition*. Harvard University Press, 1994 (Cited on pages 185, 189).

[125] Donald Norman. *The psychology of everyday things*. Vol. 5. Basic books, 1988 (Cited on page 45).

[126] Kenton O'Hara and Abigail Sellen. "A Comparison of Reading Paper and On-line Documents." In: *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. CHI '97. Atlanta, Georgia, USA: ACM, 1997, pp. 335–342. ISBN: 0-89791-802-9. DOI: 10.1145/258549.258787. URL: http://doi.acm.org/10.1145/258549.258787 (Cited on page 60).

[127] Ian Oakley and Doyoung Lee. "Interaction on the Edge: Offset Sensing for Small Devices." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. Toronto, Ontario, Canada: ACM, 2014, pp. 169–178. ISBN: 978-1-4503-2473-1. DOI: 10.1145/2556288.2557138. URL: http://doi.acm.org/10.1145/2556288.2557138 (Cited on pages 11, 14).

[128] Marianna Obrist, Sue Ann Seah, and Sriram Subramanian. "Talking About Tactile Experiences." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. Paris, France: ACM, 2013, pp. 1659–1668. ISBN: 978-1-4503-1899-0. DOI: 10.1145/2470654.2466220. URL: http://doi.acm.org/10.1145/2470654.2466220 (Cited on pages 128, 133).

[129] Masa Ogata and Michita Imai. "SkinWatch: Skin Gesture Interaction for Smart Watch." In: *Proceedings of the 6th Augmented Human International Conference on - AH '15* (2015), pp. 21–24. DOI: 10.1145/2735711.2735830 (Cited on page 14).

[130] Shogo Okamoto, Hikaru Nagano, and Yoji Yamada. "Psychophysical dimensions of tactile perception of textures." In: *IEEE Transactions on Haptics* 6.1 (2013), pp. 81–93. ISSN: 19391412. DOI: 10.1109/ToH.2012.32. URL: http://www.ncbi.nlm.nih.gov/pubmed/24808270 (Cited on pages 77, 84, 85, 122).

[131] *Omega Engineering. Strain Gages.* 2015. http://www.omega.ca/guides/straingages.html (Cited on page 50).

[132] Anke van Oosterhout, Majken Kirkegård Rasmussen, Eve Hoggan, and Miguel Bruns. "Knobology 2.0: Giving Shape to the Haptic Force Feedback of Interactive Knobs." In: *The 31st Annual ACM Symposium on User Interface Software and Technology Adjunct Proceedings*. UIST '18 Adjunct. Berlin, Germany: ACM, 2018, pp. 197–199. ISBN: 978-1-4503-5949-8. DOI: 10.1145/3266037.3271649. URL: http://doi.acm.org/10.1145/3266037.3271649 (Cited on page 188).

[133] Antti Oulasvirta, Sunjun Kim, and Byungjoo Lee. "Neuromechanics of a Button Press." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: ACM, 2018, 508:1–508:13. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3174082. URL: http://doi.acm.org/10.1145/3173574.3174082 (Cited on pages 188, 197).

[134] Abdenbi Mohand Ousaid, Guillaume Millet, Sinan Haliyo, Stéphane Régnier, and Vincent Hayward. "Feeling what an insect feels." In: *PLoS ONE* 9.10 (2014), e108895. ISSN: 19326203. DOI: 10.1371/journal.pone.0108895. URL: http://www.isir.upmc.fr/files/2014ACLI3202.pdf (Cited on pages 2, 175).

[135] Claudio Pacchierotti, Stephen Sinclair, Massimiliano Solazzi, Antonio Frisoli, Vincent Hayward, and Domenico Prattichizzo. "Wearable Haptic Systems for the Fingertip and the Hand: Taxonomy, Review, and Perspectives." In: *IEEE Transactions on Haptics* 10.4 (2017), pp. 1–1. ISSN: 1939-1412. DOI: 10.1109/TOH.2017.2689006. URL: http://ieeexplore.ieee.org/document/7922602/ (Cited on page 209).

[136] Daniel Pargman and Peter Jakobsson. "Five Perspectives on Computer Game History." In: *Interactions* 14.6 (Nov. 2007), pp. 26–29. ISSN: 1072-5520. DOI: 10.1145/1300655.1300674. URL: http://doi.acm.org/10.1145/1300655.1300674 (Cited on pages 78, 127).

[137] Simon T Perrault, Eric Lecolinet, James Eagan, and Yves Guiard. "WatchIt : Simple Gestures and Eyes-free Interaction for Wristwatches and Bracelets." In: *Proc. CHI*. 2013, pp. 1451–1460. ISBN: 9781450318990 (Cited on page 14).

[138] Claire Petitmengin-Peugeot. "The intuitive experience." In: *Journal of Consciousness Studies* 6.2-3 (1999), pp. 43–77 (Cited on pages 128, 133).

[139] Claire Petitmengin. "Describing one's subjective experience in the second person: An interview method for the science of consciousness." In: *Phenomenology and the Cognitive Sciences* 5.3 (2006), pp. 229–269. ISSN: 1572-8676. DOI: 10.1007/s11097-006-9022-2. URL: https://doi.org/10.1007/s11097-006-9022-2 (Cited on pages 128, 133).

[140] Henning Pohl and Kasper Hornbæk. "ElectricItch: Skin Irritation As a Feedback Modality." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Berlin, Germany: ACM, 2018, pp. 765–778. ISBN: 978-1-4503-5948-1. DOI: 10.1145/3242587.3242647. URL: http://doi.acm.org/10.1145/3242587.3242647 (Cited on page 5).

[141] Ivan Poupyrev, Shigeaki Maruyama, and Jun Rekimoto. "Ambient Touch: Designing Tactile Interfaces for Handheld Devices." In: *Proc. Control 2002* (2002), pp. 51–60. URL: http://www.csl.sony.co.jp/ILhttp://portal.acm.org/citation.cfm?id=571985.571993 (Cited on page 78).

[142] Ivan Poupyrev, Makoto Okabe, and Shigeaki Maruyama. "Haptic Feedback for Pen Computing: Directions and Strategies." In: *CHI '04 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '04. Vienna, Austria: ACM, 2004, pp. 1309–1312. ISBN: 1-58113-703-6. DOI: 10.1145/985921.986051. URL: http://doi.acm.org/10.1145/985921.986051 (Cited on pages 78, 128).

[143] Majken K. Rasmussen, Esben W. Pedersen, Marianne G. Petersen, and Kasper Hornbæk. "Shape-Changing Interfaces: A Review of the Design Space and Open Research Questions." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. 2012, pp. 735–744. ISBN: 9781450310154. DOI: 10.1145/2207676.2207781. URL: https://dl.acm.org/citation.cfm?doid=2207676.2207781 (Cited on page 95).

[144] Simon Robinson, Parisa Eslambolchilar, and Matt Jones. "Sweep-Shake: Finding Digital Resources in Physical Environments." In: *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services*. MobileHCI '09. Bonn, Germany: ACM, 2009, 12:1–12:10. ISBN: 978-1-60558-281-8. DOI: 10.1145/1613858.1613874. URL: http://doi.acm.org/10.1145/1613858.1613874 (Cited on page 144).

[145] Kevin Roitero, Eddy Maddalena, and Stefano Mizzaro. "On Fine-Grained Relevance Scales." In: (2018), pp. 675–684 (Cited on page 98).

[146] Joseph M. Romano and Katherine J. Kuchenbecker. "Creating Realistic Virtual Textures from Contact Acceleration Data." In: *EEE Trans. Haptics* 5.2 (Jan. 2012), pp. 109–119. ISSN: 1939-1412. DOI: 10.1109/TOH.2011.38. URL: http://dx.doi.org/10.1109/TOH.2011.38 (Cited on pages 2, 76, 79, 126, 127, 129, 132, 175, 176).

[147] David A. Rosenbaum. *Knowing Hands*. Cambridge: Cambridge University Press, 2017. ISBN: 9781316148525. DOI: 10.1017/9781316148525. URL: http://ebooks.cambridge.org/ref/id/CBO9781316148525 (Cited on pages 173, 182, 184).

[148] Joseph D. Rutledge and Ted Selker. "Force-to-motion Functions for Pointing." In: *Proceedings of the IFIP TC13 Third Interational Conference on Human-Computer Interaction*. INTERACT '90. Amsterdam, The Netherlands, The Netherlands: North-Holland Publishing Co., 1990, pp. 701–706. ISBN: 0-444-88817-9. URL: http://dl.acm.org/citation.cfm?id=647402.725310 (Cited on page 47).

[149] M Sato. *Response of Pacinian Corpuscles to Sinusoidal Vibrations*. Tech. rep. 1961, pp. 391–409. URL: https://physoc.onlinelibrary.wiley.com/doi/pdf/10.1113/jphysiol.1961.sp006817 (Cited on pages 35, 37, 110, 164).

[150] Jeff Sauro, Joseph S Dumas, and Yarmouth Port. "Comparison of Three One-Question , Post-Task Usability Questionnaires." In: (2009), pp. 1599–1608 (Cited on pages 98, 99).

[151] Robert F. Schmidt. *Fundamentals of Sensory Physiology*. Vol. 3. Berlin Heidelberg New York Tokyo, 1986, pp. 25–33. ISBN: 9788578110796. DOI: 10 . 1017 / CBO9781107415324 . 004. arXiv: arXiv : 1011 . 1669v3. URL: BusinessSourcePremierdatabase. (Cited on page 25).

[152] Stefan Schneegass and Alexandra Voit. "GestureSleeve: using touch sensitive fabrics for gestural input on the forearm for controlling smartwatches." In: *Proceedings of the 2016 ACM International Symposium on Wearable Computers - ISWC '16* (2016), pp. 108–115. DOI: 10.1145/2971763.2971797. URL: https://dl.acm.org/citation.cfm?doid=2971763.2971797 (Cited on page 14).

[153] Marian Schneider, Valentin G Kemper, Thomas C Emmerling, Federico De Martino, and Rainer Goebel. "Columnar clusters in the human motion complex reflect consciously perceived motion axis." In: *Proceedings of the National Academy of Sciences of the United States of America* 116.11 (2019), pp. 5096–5101. ISSN: 1091-6490. DOI: 10.1073/pnas.1814504116. URL: http://www.ncbi.nlm.nih.gov/pubmed/30808809 (Cited on page 184).

[154] Samuel B. Schorr and Allison M. Okamura. "Fingertip Tactile Devices for Virtual Object Manipulation and Exploration." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: ACM, 2017, pp. 3115–3119. ISBN: 978-1-4503-4655-9. DOI: 10.1145/3025453.3025744. URL: http://doi.acm.org/10.1145/3025453.3025744 (Cited on page 128).

[155] Carsten Schwesig, Ivan Poupyrev, and Eijiro Mori. "Gummi: A Bendable Computer." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '04. Vienna, Austria: ACM, 2004, pp. 263–270. ISBN: 1-58113-702-8. DOI: 10.1145/985692.985726. URL: http://doi.acm.org/10.1145/985692.985726 (Cited on page 47).

[156] Brandon Shrewsbury. "Providing Haptic Feedback Using the Kinect." In: *Journal of Rehabilitation Research* (2011), pp. 321–322. DOI: 10.1145/2049536.2049628. URL: https://dl.acm.org/citation.cfm?doid=2049536.2049628 (Cited on page 78).

[157] Advaith Siddharthan and Napoleon Katsos. "Offline Sentence Processing Measures for testing Readability with Users." In: Pitr (2012), pp. 17–24 (Cited on pages 98, 99, 101, 105).

[158] Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. "AIREAL: Interactive Tactile Experiences in Free Air." In: *ACM Trans. Graph.* 32.4 (July 2013), 134:1–134:10. ISSN: 0730-0301. DOI: 10.1145/2461912.2462007. URL: http://doi.acm.org/10.1145/2461912.2462007 (Cited on pages 15, 128).

[159] Daniel Spelmezan, Deepak Ranjan Sahoo, and Sriram Subramanian. "Sparkle: hover feedback with touchable electric arcs." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM. 2017, pp. 3705–3717 (Cited on pages 11, 12, 15).

[160] Srinath Sridhar, Anders Markussen, Antti Oulasvirta, Christian Theobalt, and Sebastian Boring. "WatchSense: On- and Above-Skin Input Sensing through aWearable Depth Sensor Srinath." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17*. New York, New York, USA: ACM Press, 2017, pp. 3891–3902. ISBN: 9781450346559. DOI: 10.1145/3025453.3026005. URL: http://dl.acm.org/citation.cfm?doid=3025453.3026005 (Cited on pages 11, 14).

[161] John K. Stevens, Robert C. Emerson, George L. Gerstein, Tamas Kallos, Gordon R. Neufeld, Charles W. Nichols, and Alan C. Rosenquist. "Paralysis of the awake human: Visual perceptions." In: *Vision Research* 16.1 (1976), 93–IN9. ISSN: 0042-6989. DOI: 10.1016/0042-6989(76)90082-1. URL: https://www.sciencedirect.com/science/article/pii/0042698976900821?via{\%}3Dihub (Cited on page 193).

[162] S. S. Stevens. "The Measurement of Loudness." In: *The Journal of the Acoustical Society of America* 27.5 (1955), pp. 815–829. ISSN: 0001-4966. DOI: 10.1121/1.1908048. URL: http://asa.scitation.org/doi/10.1121/1.1908048 (Cited on pages 76, 82, 85, 97).

[163] Paul Strohmeier. "3.18 Designing for and leveraging Active Perception." In: *On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces* (), p. 97 (Cited on page 160).

[164] Paul Strohmeier. "Exploring Bodies, Mediation and Points of View Using a Robotic Avatar." In: *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '16. Eindhoven, Netherlands: ACM, 2016, pp. 663–668. ISBN: 978-1-4503-3582-9. DOI: 10.1145/2839462.2856343. URL: http://doi.acm.org/10.1145/2839462.2856343 (Cited on pages v, 147).

[165] Paul Strohmeier. "Coupling Motion and Perception in Body Based UI." In: *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '17. Yokohama, Japan: ACM, 2017, pp. 697–701. ISBN: 978-1-4503-4676-4. DOI: 10.1145/3024969.3025038. URL: http://doi.acm.org/10.1145/3024969.3025038 (Cited on page 147).

[166] Paul Strohmeier, Sebastian Boring, and Kasper Hornbæk. "From Pulse Trains to "Coloring with Vibrations" : Motion Mappings for Mid-Air Haptic Textures." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18* (2018). DOI: 10.1145/3173574.3173639. URL: https://doi.org/10.1145/3173574.3173639 (Cited on pages 149, 159, 175, 190, 191).

[167] Paul Strohmeier, Jesse Burstyn, and Roel Vertegaal. "Effects of Display Sizes on a Scrolling Task using a Cylindrical Smartwatch." In: *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI '15* (2015), pp. 846–853. DOI: 10.1145/2786567.2793710. URL: http://dl.acm.org/citation.cfm?doid=2786567.2793710 (Cited on page 14).

[168] Paul Strohmeier and Cedric Honnet. "To improve haptic experiences in VR and AR, we must first improve tracking." In: *CHI: Conference on Human Factors in Computing Systems*. 2018 (Cited on page 176).

[169] Paul Strohmeier, Cedric Honnet, and Samppa Von Cyborg. "Developing an Ecosystem for Interactive Electronic Implants." In: *Proc. Living Machines 2016* (2016). DOI: 10.1007/978-3-319-42417-0 (Cited on page v).

[170] Paul Strohmeier and Kasper Hornbæk. "Generating Haptic Textures with a Vibrotactile Actuator." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: ACM, 2017. ISBN: 978-1-4503-4655-9. DOI: 10.1145/3025453.3025812. URL: http://doi.

acm.org/10.1145/3025453.3025812 (Cited on pages 23, 97–99, 101, 103, 105, 111, 115, 127–130, 132, 146, 174–176, 190).

[171] Paul Strohmeier, Jesse Burstyn, Juan Pablo Carrascal, Vincent Levesque, and Roel Vertegaal. "ReFlex: A Flexible Smartphone with Active Haptic Feedback for Bend Input." In: *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '16. Eindhoven, Netherlands: ACM, 2016. ISBN: 978-1-4503-3582-9. DOI: 10.1145/2839462.2839494. URL: http://doi.acm.org/10.1145/2839462.2839494 (Cited on pages 76, 78, 79, 126–128, 130, 175, 176, 190).

[172] Paul Strohmeier, Jarrod Knibbe, Sebastian Boring, and Kasper Hornbæk. "zPatch: Hybrid Resistive/Capacitive eTextile Input." In: *Tangible and Embedded Interaction*. 2018, pp. 188–198 (Cited on page 6).

[173] Paul Strohmeier, Victor Håkansson, Cedric Honnet, Kasper Hornbæk, and Daniel Ashbrook. "Optimizing Pressure Matrices: Interdigitation and Interpolation Methods for Continuous Position Input." In: *ACM International Conference on Tangible, Embedded and Embodied Interaction (TEI)*. 2019 (Cited on page 6).

[174] Yuriko Suzuki and Minoru Kobayashi. "Air jet driven force feedback in virtual reality." In: *IEEE computer graphics and applications* 25.1 (2005), pp. 44–47 (Cited on page 15).

[175] Dag Svanæs. "Interaction Design for and with the Lived Body: Some Implications of Merleau-ponty's Phenomenology." In: *ACM Trans. Comput.-Hum. Interact.* 20.1 (Apr. 2013), 8:1–8:30. ISSN: 1073-0516. DOI: 10.1145/2442106.2442114. URL: http://doi.acm.org/10.1145/2442106.2442114 (Cited on pages 147, 183, 195).

[176] F.M. Swain. "Composition of marsh gases in the central and eastern United States." In: *Applied Geochemistry* 1.2 (1986), pp. 301 –305. ISSN: 0883-2927. DOI: https://doi.org/10.1016/0883-2927(86)90014-4. URL: http://www.sciencedirect.com/science/article/pii/0883292786900144 (Cited on page 198).

[177] *Tactile Labs. Hapuator Mark II.* 2012. http://www.tactilelabs.com/products/haptics/haptuator-mark-ii-v2/ (Cited on page 51).

[178] W H Talbot, I Darian-Smith, H H Kornhuber, and V B Mountcastle. "The sense of flutter-vibration: comparison of the human capacity with response patterns of mechanoreceptive afferents from the monkey hand." In: *Journal of Neurophysiology* 31.2

(1967), pp. 301–334. ISSN: 0022-3077. DOI: 10.1152/jn.1968.31.2.301 (Cited on page 35).

[179] John Tiab and Kasper Hornbæk. "Understanding Affordance, System State, and Feedback in Shape-Changing Buttons." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: ACM, 2016, pp. 2752–2763. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858350. URL: http://doi.acm.org/10.1145/2858036.2858350 (Cited on page 188).

[180] D. Todorovic. "Gestalt principles." In: *Scholarpedia* 3.12 (2008). revision #91314, p. 5345. DOI: 10.4249/scholarpedia.5345 (Cited on page 183).

[181] T O M Troscianko, Caterina Ripamonti, and David J Tolhurst. "Evaluation of a Multiscale Color Model for Visual Difference Prediction." In: 3.3 (2006), pp. 155–178 (Cited on pages 98, 99).

[182] Brygg Ullmer and Hiroshi Ishii. "The metaDESK: models and prototypes for tangible user interfaces." In: *Proceedings of the 10th annual ACM symposium on User interface software and technology - UIST '97* (1997), pp. 223–232. DOI: 10.1145/263407.263551. URL: https://dl.acm.org/citation.cfm?doid=263407.263551 (Cited on page 95).

[183] Camila Valenzuela Moguillansky, J.Kevin O'Regan, and Claire Petitmengin. "Exploring the subjective experience of the "rubber hand" illusion." In: *Frontiers in Human Neuroscience* 7 (2013), p. 659. ISSN: 1662-5161. DOI: 10.3389/fnhum.2013.00659. URL: https://www.frontiersin.org/article/10.3389/fnhum.2013.00659 (Cited on page 133).

[184] Jelle Van Dijk. "Designing for Embodied Being-in-the-World: A Critical Analysis of the Concept of Embodiment in the Design of Hybrids." In: *Multimodal Technologies and Interaction* 2.1 (2018). ISSN: 2414-4088. DOI: 10.3390/mti2010007. URL: http://www.mdpi.com/2414-4088/2/1/7 (Cited on page 207).

[185] Radu-Daniel Vatavu, Annette Mossel, and Christian Schönauer. "Digital Vibrons: Understanding Users' Perceptions of Interacting with Invisible, Zero-weight Matter." In: *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*. MobileHCI '16. Florence, Italy: ACM, 2016, pp. 217–226. ISBN: 978-1-4503-4408-1. DOI: 10.1145/2935334.2935364. URL: http://doi.acm.org/10.1145/2935334.2935364 (Cited on page 11).

[186] Peter-Paul Verbeek. *What things do: Philosophical reflections on technology, agency, and design*. Penn State Press, 2005 (Cited on pages 3, 147, 173, 177, 182).

[187] Ronald T Verrillo. "Effect of Contactor Area on the Vibrotactile Threshold." In: *The Journal of the Acoustical Society of America* 35.12 (1962). ISSN: 0001-4966. DOI: 10.1121/1.1918868. URL: https://doi.org/10.1121/1.1918868 (Cited on pages 35–37, 41, 42, 164).

[188] Ronald T. Verrillo. "Vibrotactile sensitivity and the frequency response of the Pacinian corpuscle." In: *Psychonomic Science* 4.1 (2014), pp. 135–136. ISSN: 0033-3131. DOI: 10.3758/bf03342215. URL: https://link.springer.com/content/pdf/10.3758{\%}2FBF03342215.pdf (Cited on pages 35, 37, 110, 164).

[189] Mélodie Vidal, Andreas Bulling, and Hans Gellersen. "Pursuits: Spontaneous Interaction with Displays Based on Smooth Pursuit Eye Movement and Moving Targets." In: *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. UbiComp '13. Zurich, Switzerland: ACM, 2013, pp. 439–448. ISBN: 978-1-4503-1770-2. DOI: 10.1145/2493432.2493477. URL: http://doi.acm.org/10.1145/2493432.2493477 (Cited on page 186).

[190] Johan Wagemans, James H Elder, Michael Kubovy, Stephen E Palmer, Mary A Peterson, Manish Singh, and Rüdiger von der Heydt. "A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization." In: *Psychological bulletin* 138.6 (2012), pp. 1172–217. ISSN: 1939-1455. DOI: 10.1037/a0029333. URL: http://www.ncbi.nlm.nih.gov/pubmed/22845751http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3482144 (Cited on page 183).

[191] Bruce N Walker and Lisa M Mauney. "Universal Design of Auditory Graphs : A Comparison of Sonification Mappings for Visually Impaired and Sighted Listeners." In: 2.3 (2010), pp. 1–16. DOI: 10.1145/1714458.1714459.http (Cited on pages 98, 99, 101, 111).

[192] Julie M Walker, Heather Culbertson, Michael Raitor, and Allison M Okamura. "Haptic orientation guidance using two parallel double-gimbal control moment gyroscopes." In: *IEEE transactions on haptics* 11.2 (2018), pp. 267–278 (Cited on page 127).

[193] Steven Wall and William Harwin. "Quantification of the effects of haptic feedback during a motor skills task in a simulated environment." In: *Proc. Second PHANToM users research symposium*. 2000 (Cited on page 128).

[194] Malte Weiss, Chat Wacharamanotham, Simon Voelker, and Jan Borchers. "FingerFlux: near-surface haptic feedback on tabletops." In: *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM. 2011, pp. 615–620 (Cited on pages 15, 31).

[195] David A. Whetten. "What Constitutes a Theoretical Contribution?" In: *Academy of Management Review* 14.4 (2011), pp. 490–495. ISSN: 0363-7425. DOI: 10.5465/amr.1989.4308371. arXiv: arXiv:1011.1669v3. URL: http://journals.aom.org/doi/10.5465/amr.1989.4308371 (Cited on page 192).

[196] Peter Williams and Daniel Overholt. "bEADS Extended Actuated Digital Shaker." In: *NIME 2017 Proceedings of the International Conference on New Interfaces for Musical Expression*. 2017, pp. 13–18. URL: http://homes.create.aau.dk/dano/nime17/papers/0003/paper0003.pdf (Cited on page 190).

[197] Graham Wilson, Martin Halvey, Stephen A. Brewster, and Stephen A. Hughes. "Some Like It Hot: Thermal Feedback for Mobile Devices." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pp. 2555–2564. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1979316. URL: http://doi.acm.org/10.1145/1978942.1979316 (Cited on page 128).

[198] Graham Wilson, Thomas Carter, Sriram Subramanian, and Stephen A. Brewster. "Perception of Ultrasonic Haptic Feedback on the Hand: Localisation and Apparent Motion." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. Toronto, Ontario, Canada: ACM, 2014, pp. 1133–1142. ISBN: 978-1-4503-2473-1. DOI: 10.1145/2556288.2557033. URL: http://doi.acm.org/10.1145/2556288.2557033 (Cited on page 12).

[199] Laura Winfield, John Glassmire, J Edward Colgate, and Michael Peshkin. "T-PaD: Tactile Pattern Display through Variable Friction Reduction." In: *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)* (Cited on page 78).

[200] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. "The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pp. 143–146. ISBN: 978-1-4503-0228-9. DOI: 10.1145/1978942.1978963. URL: http://doi.acm.org/10.1145/1978942.1978963 (Cited on page 99).

[201] Matthew Wright and Adrian Freed. "Open Sound Control: A New Protocol for Communicating with Sound Synthesizers." In: *Proc. ICMC 1997*. International Computer Music Association, 1997, pp. 101–104. URL: http://cnmat.berkeley.edu/publications/open-sound-control-new-protocol-communicating-sound-synthesizers (Cited on page 82).

[202] Chien-Te Wu, Sébastien M. Crouzet, Simon J. Thorpe, and Michele Fabre-Thorpe. "At 120 msec You Can Spot the Animal but You Don't Yet Know It's a Dog." In: *Journal of Cognitive Neuroscience* 27.1 (2015). PMID: 25208739, pp. 141–149. DOI: 10.1162/jocn_a_00701. eprint: https://doi.org/10.1162/jocn_a_00701. URL: https://doi.org/10.1162/jocn_a_00701 (Cited on page 186).

[203] Juan Wu, Na Li, Wei Liu, Guangming Song, and Jun Zhang. "Experimental Study on the Perception Characteristics of Haptic Texture by Multidimensional Scaling." In: *IEEE Transactions on Haptics* 8.4 (2015), pp. 410–420. ISSN: 19391412. DOI: 10.1109/TOH.2015.2438866. URL: http://www.ncbi.nlm.nih.gov/pubmed/26054074 (Cited on page 77).

[204] Danfei Xu, Gerald E Loeb, and Jeremy A Fishel. "Tactile identification of objects using Bayesian exploration." In: *IEEE International Conference on Robotics and Automation (ICRA)* (2013) (Cited on page 79).

[205] Hsin-Yun Yao and Vincent Hayward. "Design and analysis of a recoil-type vibrotactile transducer." In: *The Journal of the Acoustical Society of America* 128.2 (2010), pp. 619–627 (Cited on pages 78, 126, 128).

[206] Hsin-yun Yao and Vincent Hayward. "An experiment on length perception with a virtual rolling stone." In: *Proceedings of Eurohaptics*. 2006, pp. 325–330 (Cited on pages 46, 48, 76, 78, 79, 128, 175).

[207] Alfred L. Yarbus. *Eye Movement and Vision*. Ed. by Lorrin A. Riggs. New York: Plenum Press, 1967 (Cited on pages 175, 193).

[208] Takashi Yoshioka, James C Craig, Graham C Beck, and Steven S Hsiao. "Perceptual constancy of texture roughness in the tactile system." In: *J Neurosci* 31.48 (2011), pp. 17603–17611. ISSN: 1529-2401. DOI: 10.1523/JNEUROSCI.3907-11.2011. URL: http://dx.doi.org/10.1523/JNEUROSCI.3907-11.2011 (Cited on page 77).

[209] Shumin Zhai. *Human performance in six degree of freedom input control*. University of Toronto, 1996 (Cited on pages 47, 51, 60, 165, 167).

[210] Shumin Zhai. "Human performance." In: *Haptic Input*. Ed. by Bill Buxton. 2009, pp. 1–28 (Cited on page 165).

[211] Siyan Zhao, Ali Israr, and Roberta Klatzky. "Intermanual apparent tactile motion on handheld tablets." In: *2015 IEEE World Haptics Conference (WHC)*. IEEE. 2015, pp. 241–247 (Cited on pages 78, 128).

[212] Siyan Zhao, Oliver Schneider, Roberta Klatzky, Jill Lehman, and Ali Israr. "FeelCraft: Crafting Tactile Experiences for Media Using a Feel Effect Library." In: *Proceedings of the Adjunct Publication of the 27th Annual ACM Symposium on User Interface Software and Technology*. UIST'14 Adjunct. Honolulu, Hawaii, USA: ACM, 2014, pp. 51–52. ISBN: 978-1-4503-3068-8. DOI: 10.1145/2658779.2659109. URL: http://doi.acm.org/10.1145/2658779.2659109 (Cited on pages 78, 128).

[213] Siyan Zhao, Zachary Schwemler, Adam Fritz, and Ali Israr. "Stereo Haptics: Designing Haptic Interactions Using Audio Tools." In: *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '16. Eindhoven, Netherlands: ACM, 2016, pp. 778–781. ISBN: 978-1-4503-3582-9. DOI: 10.1145/2839462.2854120. URL: http://doi.acm.org/10.1145/2839462.2854120 (Cited on pages 78, 128).

[214] Junhan Zhou, Yang Zhang, Gierad Laput, and Chris Harrison. "AuraSense: enabling expressive around-smartwatch interactions with electric field sensing." In: *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM. 2016, pp. 81–86 (Cited on page 14).

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both LaTeX and LYX:
https://bitbucket.org/amiede/classicthesis/

Updated Version incorporating feedback provided by Dan Vogel (July 2019).